

**Dell Networking Configuration Guide for the
MXL 10/40GbE Switch I/O Module
9.8(0.0)**



Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your computer.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Copyright © 2015 Dell Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. Dell™ and the Dell logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

2015 - 05

Rev. A00

Contents

1 About this Guide.....	33
Audience.....	33
Conventions.....	33
Information Symbols.....	34
Related Documents.....	34
2 Configuration Fundamentals.....	35
Accessing the Command Line.....	35
CLI Modes.....	35
Navigating CLI Modes.....	37
The do Command.....	40
Undoing Commands.....	40
Obtaining Help.....	41
Entering and Editing Commands.....	42
Command History.....	43
Filtering show Command Outputs.....	43
Multiple Users in Configuration Mode.....	44
3 Getting Started.....	46
Console Access.....	47
Serial Console.....	47
External Serial Port with a USB Connector.....	49
Accessing the CLI Interface and Running Scripts Using SSH.....	49
Entering CLI commands Using an SSH Connection.....	49
Executing Local CLI Scripts Using an SSH Connection.....	49
Boot Process.....	50
Default Configuration.....	52
Configuring a Host Name.....	52
Configuring a Host Name.....	52
Accessing the System Remotely.....	53
Accessing the MXL Switch Remotely.....	53
Configure the Management Port IP Address.....	53
Configure a Management Route.....	53
Configuring a Username and Password.....	54
Configuring the Enable Password.....	54
Configuration File Management.....	55
Copy Files to and from the System.....	55
Save the Running-Configuration.....	56

Viewing Files.....	57
Managing the File System.....	58
View the Command History.....	59
Using HTTP for File Transfers.....	60
Upgrading and Downgrading the Dell Networking OS.....	60
Using Hashes to Validate Software Images.....	60
4 Management.....	62
Configuring Privilege Levels.....	62
Creating a Custom Privilege Level.....	62
Customizing a Privilege Level.....	63
Applying a Privilege Level to a Username.....	64
Applying a Privilege Level to a Terminal Line.....	65
Configuring Logging.....	65
Audit and Security Logs.....	65
Configuring Logging Format	67
Setting Up a Secure Connection to a Syslog Server.....	68
Display the Logging Buffer and the Logging Configuration.....	69
Log Messages in the Internal Buffer.....	70
Configuration Task List for System Log Management.....	70
Disabling System Logging.....	70
Sending System Messages to a Syslog Server.....	70
Configuring a UNIX System as a Syslog Server.....	71
Changing System Logging Settings.....	71
Display the Logging Buffer and the Logging Configuration.....	72
Configuring a UNIX Logging Facility Level.....	73
Synchronizing Log Messages.....	74
Enabling Timestamp on Syslog Messages.....	74
File Transfer Services.....	75
Configuration Task List for File Transfer Services.....	75
Enabling the FTP Server.....	75
Configuring FTP Server Parameters.....	76
Configuring FTP Client Parameters.....	76
Terminal Lines.....	77
Denying and Permitting Access to a Terminal Line.....	77
Configuring Login Authentication for Terminal Lines.....	78
Setting Time Out of EXEC Privilege Mode.....	79
Using Telnet to get to Another Network Device.....	79
Lock CONFIGURATION Mode.....	80
Viewing the Configuration Lock Status.....	80
Limit Concurrent Login Sessions.....	81
Restrictions for Limiting the Number of Concurrent Sessions.....	81

Configuring Concurrent Session Limit.....	81
Enabling the System to Clear Existing Sessions.....	82
Track Login Activity.....	83
Restrictions for Tracking Login Activity.....	83
Configuring Login Activity Tracking.....	83
Display Login Statistics.....	84
Recovering from a Forgotten Password.....	85
Recovering from a Forgotten Enable Password.....	86
Recovering from a Failed Start.....	86
5 802.1X.....	88
The Port-Authentication Process.....	90
EAP over RADIUS.....	92
Configuring 802.1X.....	92
Related Configuration Tasks.....	92
Important Points to Remember.....	93
Enabling 802.1X.....	93
Configuring Request Identity Re-Transmissions.....	95
Configuring a Quiet Period after a Failed Authentication.....	95
Forcibly Authorizing or Unauthorizing a Port.....	96
Re-Authenticating a Port.....	97
Configuring Timeouts.....	98
Configuring Dynamic VLAN Assignment with Port Authentication.....	99
Guest and Authentication-Fail VLANs.....	100
Configuring a Guest VLAN.....	101
Configuring an Authentication-Fail VLAN.....	101
6 Access Control List (ACL) VLAN Groups and Content Addressable Memory (CAM).....	103
Optimizing CAM Utilization During the Attachment of ACLs to VLANs.....	103
Guidelines for Configuring ACL VLAN groups.....	104
Configuring ACL VLAN Groups and Configuring FP Blocks for VLAN Parameters.....	105
Configuring ACL VLAN Groups.....	105
Configuring FP Blocks for VLAN Parameters.....	106
Viewing CAM Usage.....	107
Allocating FP Blocks for VLAN Processes.....	108
7 Access Control Lists (ACLs).....	110
IP Access Control Lists (ACLs).....	110
Implementing ACL on the Dell Networking OS.....	111
ACLs and VLANs.....	111
ACL Optimization.....	111

Determine the Order in which ACLs are Used to Classify Traffic.....	111
Example of the order Keyword to Determine ACL Sequence.....	112
IP Fragment Handling.....	112
IP Fragments ACL Examples.....	113
Layer 4 ACL Rules Examples.....	113
Configure a Standard IP ACL.....	114
Configuring a Standard IP ACL Filter.....	115
Configure an Extended IP ACL.....	116
Configuring Filters with a Sequence Number.....	116
Configuring Filters Without a Sequence Number.....	117
Established Flag.....	118
Configure Layer 2 and Layer 3 ACLs.....	118
Assign an IP ACL to an Interface.....	119
Applying an IP ACL.....	119
Counting ACL Hits.....	120
Configure Ingress ACLs.....	120
Configure Egress ACLs.....	121
Applying Egress Layer 3 ACLs (Control-Plane).....	121
IP Prefix Lists.....	122
Implementation Information.....	123
Configuration Task List for Prefix Lists.....	123
Creating a Prefix List.....	123
Creating a Prefix List Without a Sequence Number.....	124
Viewing Prefix Lists.....	125
Applying a Prefix List for Route Redistribution.....	125
Applying a Filter to a Prefix List (OSPF).....	126
ACL Resequencing.....	127
Resequencing an ACL or Prefix List.....	127
Route Maps.....	129
Implementation Information.....	129
Important Points to Remember.....	129
Configuration Task List for Route Maps.....	130
Creating a Route Map.....	130
Configure Route Map Filters.....	131
Configuring Match Routes.....	132
Configuring Set Conditions.....	133
Configure a Route Map for Route Redistribution.....	134
Configure a Route Map for Route Tagging.....	135
Continue Clause.....	135
Logging of ACL Processes.....	136
Guidelines for Configuring ACL Logging.....	137
Configuring ACL Logging.....	137

Flow-Based Monitoring Support for ACLs.....	138
Behavior of Flow-Based Monitoring.....	138
Enabling Flow-Based Monitoring.....	140
8 Bidirectional Forwarding Detection (BFD).....	142
How BFD Works.....	142
BFD Packet Format.....	143
BFD Sessions.....	145
BFD Three-Way Handshake.....	145
Session State Changes.....	146
Important Points to Remember.....	147
Configure BFD.....	148
Configure BFD for Physical Ports.....	148
Enabling BFD Globally.....	148
Establishing a Session on Physical Ports.....	149
Changing Physical Port Session Parameters.....	150
Disabling and Re-Enabling BFD.....	151
Configure BFD for Static Routes.....	152
Related Configuration Tasks.....	152
Establishing Sessions for Static Routes.....	152
Changing Static Route Session Parameters.....	153
Disabling BFD for Static Routes.....	153
Configure BFD for OSPF.....	154
Related Configuration Tasks.....	154
Establishing Sessions with OSPF Neighbors.....	155
Changing OSPF Session Parameters.....	156
Disabling BFD for OSPF.....	157
Configure BFD for OSPFv3.....	157
Related Configuration Tasks.....	157
Establishing Sessions with OSPFv3 Neighbors.....	157
Changing OSPFv3 Session Parameters.....	158
Disabling BFD for OSPFv3.....	158
Configure BFD for BGP.....	159
Prerequisites.....	159
Establishing Sessions with BGP Neighbors.....	159
Disabling BFD for BGP.....	161
Use BFD in a BGP Peer Group.....	162
Displaying BFD for BGP Information.....	162
Configure BFD for VRRP.....	166
Related Configuration Tasks.....	167
Establishing Sessions with All VRRP Neighbors.....	167
Establishing VRRP Sessions on VRRP Neighbors.....	167

Changing VRRP Session Parameters.....	168
Disabling BFD for VRRP.....	169
Configure BFD for VLANs.....	169
Related Configuration Task.....	170
Establish Sessions with VLAN Neighbors.....	170
Changing VLAN Session Parameters.....	171
Disabling BFD for VLANs.....	171
Configure BFD for Port-Channels.....	171
Related Configuration Tasks.....	172
Establish Sessions on Port-Channels.....	172
Changing Physical Port Session Parameters.....	173
Disabling BFD for Port-Channels.....	173
Configuring Protocol Liveness.....	173
Troubleshooting BFD.....	174
9 Border Gateway Protocol IPv4 (BGPv4).....	175
Autonomous Systems (AS).....	175
Sessions and Peers.....	177
Establish a Session.....	178
Route Reflectors.....	179
Communities.....	179
BGP Attributes.....	180
Best Path Selection Criteria.....	180
Weight.....	182
Local Preference.....	182
Multi-Exit Discriminators (MEDs).....	183
Origin.....	184
AS Path.....	185
Next Hop.....	185
Multiprotocol BGP.....	186
Implement BGP with the Dell Networking OS.....	186
Additional Path (Add-Path) Support.....	186
Advertise IGP Cost as MED for Redistributed Routes.....	186
Ignore Router-ID for Some Best-Path Calculations.....	187
Four-Byte AS Numbers.....	187
AS4 Number Representation.....	188
AS Number Migration.....	190
BGP4 Management Information Base (MIB).....	191
Important Points to Remember.....	191
Configuration Information.....	192
BGP Configuration.....	193
Enabling BGP.....	194

Enabling MBGP Configurations.....	227
BGP Regular Expression Optimization.....	228
Debugging BGP.....	228
Storing Last and Bad PDUs.....	229
PDU Counters.....	230
Sample Configurations.....	230
10 Content Addressable Memory (CAM).....	240
CAM Allocation.....	240
Test CAM Usage.....	241
View CAM-ACL Settings.....	242
CAM Optimization.....	242
11 Control Plane Policing (CoPP).....	243
Configure Control Plane Policing.....	244
Configuring CoPP for Protocols.....	245
Configuring CoPP for CPU Queues.....	247
Show Commands.....	248
12 Data Center Bridging (DCB).....	250
Ethernet Enhancements in Data Center Bridging.....	250
Priority-Based Flow Control.....	251
Enhanced Transmission Selection.....	252
Data Center Bridging Exchange Protocol (DCBx).....	254
Data Center Bridging in a Traffic Flow.....	254
Enabling Data Center Bridging.....	255
Configuring DCB Maps and its Attributes.....	256
Data Center Bridging: Default Configuration.....	259
Interworking of DCB Map With DCB Buffer Threshold Settings.....	260
Configuring Priority-Based Flow Control.....	261
Configuring Lossless Queues.....	262
Configuring the PFC Buffer in a Switch Stack.....	263
Priority-Based Flow Control Using Dynamic Buffer Method.....	264
Configure Enhanced Transmission Selection.....	265
ETS Prerequisites and Restrictions.....	266
Creating an ETS Priority Group.....	266
ETS Operation with DCBx.....	267
Configuring Bandwidth Allocation for DCBx CIN.....	268
Hierarchical Scheduling in ETS Output Policies.....	269
Applying DCB Policies with an ETS Configuration.....	269
PFC and ETS Configuration Examples.....	270
Using PFC and ETS to Manage Data Center Traffic.....	270

Using PFC and ETS to Manage Converged Ethernet Traffic in a Switch Stack.....	272
Applying DCB Policies in a Switch Stack.....	273
Configure a DCBx Operation.....	273
DCBx Operation.....	273
DCBx Port Roles.....	274
DCB Configuration Exchange.....	276
Configuration Source Election.....	276
Propagation of DCB Information.....	276
Auto-Detection and Manual Configuration of the DCBx Version.....	277
DCBx Example.....	277
DCBx Prerequisites and Restrictions.....	278
Configuring DCBx.....	279
Verifying the DCB Configuration.....	283
QoS dot1p Traffic Classification and Queue Assignment.....	293
Configuring the Dynamic Buffer Method.....	294
13 Debugging and Diagnostics.....	296
Offline Diagnostics.....	296
Important Points to Remember.....	296
Running Offline Diagnostics.....	296
Trace Logs.....	299
Auto Save on Crash or Rollover.....	299
Using the Show Hardware Commands.....	300
Enabling Environmental Monitoring.....	301
Recognize an Over-Temperature Condition.....	303
Troubleshoot an Over-Temperature Condition.....	303
Recognize an Under-Voltage Condition.....	304
Troubleshoot an Under-Voltage Condition.....	304
Troubleshooting Packet Loss.....	305
Displaying Drop Counters.....	305
Dataplane Statistics.....	306
Display Stack Port Statistics.....	307
Displaying Stack Member Counters.....	307
Enabling Application Core Dumps.....	308
Mini Core Dumps.....	308
Enabling TCP Dumps.....	309
Enabling Buffer Statistics Tracking	310
14 Dynamic Host Configuration Protocol (DHCP).....	311
DHCP Packet Format and Options.....	311
Assign an IP Address using DHCP.....	313
Implementation Information.....	314

Configure the System to be a DHCP Server.....	315
Configuring the Server for Automatic Address Allocation.....	315
Configuration Tasks.....	316
Specifying a Default Gateway.....	317
Enabling the DHCP Server.....	317
Configure a Method of Hostname Resolution.....	318
Creating Manual Binding Entries.....	319
Debugging the DHCP Server.....	319
Using DHCP Clear Commands.....	319
Configure the System to be a Relay Agent.....	320
Configure the System to be a DHCP Client.....	322
Configuring the DHCP Client System.....	322
DHCP Client on a Management Interface.....	326
DHCP Client Operation with Other Features.....	327
Configure Secure DHCP.....	328
Option 82.....	328
DHCP Snooping.....	329
Drop DHCP Packets on Snooped VLANs Only.....	332
Dynamic ARP Inspection.....	333
Configuring Dynamic ARP Inspection.....	334
Source Address Validation.....	335
15 Equal Cost Multi-Path (ECMP).....	337
ECMP for Flow-Based Affinity.....	337
Enabling Deterministic ECMP Next Hop.....	337
Link Bundle Monitoring.....	338
Managing ECMP Group Paths.....	338
16 FC FLEXIO FPORT.....	340
FC FLEXIO FPORT.....	340
Configuring Switch Mode to FCF Port Mode.....	340
Name Server.....	341
FCoE Maps.....	342
Creating an FCoE Map.....	342
Zoning.....	344
Creating Zone and Adding Members.....	344
Creating Zone Alias and Adding Members.....	345
Creating Zonesets.....	345
Activating a Zoneset.....	346
Displaying the Fabric Parameters.....	346
17 FCoE Transit.....	349

Fibre Channel over Ethernet.....	349
Ensure Robustness in a Converged Ethernet Network.....	349
FIP Snooping on Ethernet Bridges.....	351
FIP Snooping in a Switch Stack.....	354
Using FIP Snooping.....	354
Important Points to Remember.....	354
Enabling the FCoE Transit Feature.....	355
Enable FIP Snooping on VLANs.....	355
Configure the FC-MAP Value.....	355
Configure a Port for a Bridge-to-Bridge Link.....	355
Configure a Port for a Bridge-to-FCF Link.....	356
Impact on Other Software Features.....	356
FIP Snooping Prerequisites.....	356
FIP Snooping Restrictions.....	357
Configuring FIP Snooping.....	357
Displaying FIP Snooping Information.....	358
FCoE Transit Configuration Example.....	364
18 FIPS Cryptography.....	366
Preparing the System.....	366
Enabling FIPS Mode.....	366
Generating Host-Keys.....	367
Monitoring FIPS Mode Status.....	367
Disabling FIPS Mode.....	368
19 Force10 Resilient Ring Protocol (FRRP).....	370
Protocol Overview.....	370
Ring Status.....	371
Multiple FRRP Rings.....	372
Important FRRP Points.....	373
Important FRRP Concepts.....	374
Implementing FRRP.....	375
FRRP Configuration.....	375
Creating the FRRP Group.....	376
Configuring the Control VLAN.....	376
Configuring and Adding the Member VLANs.....	377
Setting the FRRP Timers.....	379
Clearing the FRRP Counters.....	379
Viewing the FRRP Configuration.....	379
Viewing the FRRP Information.....	380
Troubleshooting FRRP.....	380
Configuration Checks.....	380

Sample Configuration and Topology.....	381
20 GARP VLAN Registration Protocol (GVRP).....	383
Important Points to Remember.....	383
Configure GVRP.....	384
Related Configuration Tasks.....	385
Enabling GVRP Globally.....	385
Enabling GVRP on a Layer 2 Interface.....	385
Configure GVRP Registration.....	386
Configure a GARP Timer.....	386
21 Internet Group Management Protocol (IGMP).....	388
IGMP Protocol Overview.....	388
IGMP Version 2.....	388
IGMP Version 3.....	390
IGMP Snooping.....	393
IGMP Snooping Implementation Information.....	393
Configuring IGMP Snooping.....	394
Enabling IGMP Immediate-Leave.....	394
Disabling Multicast Flooding.....	395
Specifying a Port as Connected to a Multicast Router.....	395
Configuring the Switch as Querier.....	395
Fast Convergence after MSTP Topology Changes.....	396
Designating a Multicast Router Interface.....	396
22 Interfaces.....	397
Basic Interface Configuration.....	397
Advanced Interface Configuration.....	397
Interface Types.....	398
View Basic Interface Information.....	398
Enabling a Physical Interface.....	400
Physical Interfaces.....	401
Configuration Task List for Physical Interfaces.....	401
Overview of Layer Modes.....	401
Configuring Layer 2 (Data Link) Mode.....	402
Configuring Layer 2 (Interface) Mode.....	402
Configuring Layer 3 (Network) Mode.....	403
Configuring Layer 3 (Interface) Mode.....	403
Management Interfaces.....	404
Configuring Management Interfaces on the MXL Switch.....	405
VLAN Interfaces.....	406
Loopback Interfaces.....	407

Null Interfaces.....	408
Port Channel Interfaces.....	408
Port Channel Definition and Standards.....	408
Port Channel Benefits.....	408
Port Channel Implementation.....	409
100/1000/10000 Mbps Interfaces in Port Channels.....	409
Configuration Tasks for Port Channel Interfaces.....	410
Creating a Port Channel.....	410
Adding a Physical Interface to a Port Channel.....	410
Reassigning an Interface to a New Port Channel.....	412
Configuring the Minimum Oper Up Links in a Port Channel.....	413
Adding or Removing a Port Channel from a VLAN.....	413
Assigning an IP Address to a Port Channel.....	414
Deleting or Disabling a Port Channel.....	414
Load Balancing through Port Channels.....	415
Load-Balancing Method.....	415
Hash Algorithm.....	416
Server Ports.....	416
Default Configuration without Start-up Config.....	417
Bulk Configuration.....	417
Interface Range.....	417
Bulk Configuration Examples.....	418
Defining Interface Range Macros.....	419
Define the Interface Range.....	420
Choosing an Interface-Range Macro.....	420
Monitoring and Maintaining Interfaces.....	420
Maintenance Using TDR.....	421
Splitting QSFP Ports to SFP+ Ports.....	422
Merging SFP+ Ports to QSFP 40G Ports.....	423
Configure the MTU Size on an Interface.....	423
Converting a QSFP or QSFP+ Port to an SFP or SFP+ Port.....	424
Important Points to Remember.....	425
Example Scenarios.....	425
Layer 2 Flow Control Using Ethernet Pause Frames.....	425
Enabling Pause Frames.....	426
Configure MTU Size on an Interface.....	427
Port-Pipes.....	428
Auto-Negotiation on Ethernet Interfaces.....	428
Setting the Speed and Duplex Mode of Ethernet Interfaces.....	428
View Advanced Interface Information.....	430
Configuring the Interface Sampling Size.....	431
Dynamic Counters.....	432

Enhanced Validation of Interface Ranges.....	434
Enhanced Control of Remote Fault Indication Processing.....	434
23 Internet Protocol Security (IPSec).....	435
Configuring IPSec	436
24 IPv4 Routing.....	437
IP Addresses.....	437
Implementation Information.....	437
Configuration Tasks for IP Addresses.....	438
IPv4 Path MTU Discovery Overview.....	441
Using the Configured Source IP Address in ICMP Messages.....	441
Configuring the ICMP Source Interface.....	442
Configuring the Duration to Establish a TCP Connection.....	442
Enabling Directed Broadcast.....	443
Resolution of Host Names.....	443
Enabling Dynamic Resolution of Host Names.....	443
Specifying the Local System Domain and a List of Domains.....	444
Configuring DNS with Traceroute.....	444
ARP.....	445
Configuration Tasks for ARP.....	446
ARP Learning via Gratuitous ARP.....	447
ARP Learning via ARP Request.....	448
Configuring ARP Retries.....	449
ICMP.....	449
Configuration Tasks for ICMP.....	449
UDP Helper.....	450
Configure UDP Helper.....	450
Important Points to Remember.....	450
Enabling UDP Helper.....	450
Configurations Using UDP Helper.....	451
UDP Helper with Broadcast-All Addresses.....	451
UDP Helper with Subnet Broadcast Addresses.....	452
UDP Helper with Configured Broadcast Addresses.....	453
UDP Helper with No Configured Broadcast Addresses.....	453
Troubleshooting UDP Helper.....	453
25 IPv6 Addressing.....	455
Protocol Overview.....	455
Extended Address Space.....	455
Stateless Autoconfiguration.....	455
IPv6 Header Fields.....	457

Version (4 bits).....	457
Traffic Class (8 bits).....	457
Flow Label (20 bits).....	457
Payload Length (16 bits).....	458
Next Header (8 bits).....	458
Hop Limit (8 bits).....	458
Source Address (128 bits).....	459
Destination Address (128 bits).....	459
Extension Header Fields.....	459
Hop-by-Hop Options Header.....	459
Addressing.....	460
Link-local Addresses.....	461
Static and Dynamic Addressing.....	461
Implementing IPv6 with the Dell Networking OS.....	461
ICMPv6.....	464
Path MTU Discovery.....	464
IPv6 Neighbor Discovery.....	465
IPv6 Neighbor Discovery of MTU Packets.....	466
Configuring the IPv6 Recursive DNS Server.....	466
Debugging IPv6 RDNSS Information Sent to the Host	467
Displaying IPv6 RDNSS Information.....	467
IPv6 Multicast.....	467
Secure Shell (SSH) Over an IPv6 Transport.....	468
Configuration Task List for IPv6.....	468
Adjusting Your CAM-Profile.....	468
Assigning an IPv6 Address to an Interface.....	469
Assigning a Static IPv6 Route.....	470
Configuring Telnet with IPv6.....	470
SNMP over IPv6.....	471
Showing IPv6 Information.....	471
Showing an IPv6 Interface.....	471
Showing IPv6 Routes.....	472
Showing the Running-Configuration for an Interface.....	473
Clearing IPv6 Routes.....	474
26 iSCSI Optimization.....	475
iSCSI Optimization Overview.....	475
Monitoring iSCSI Traffic Flows.....	477
Information Monitored in iSCSI Traffic Flows.....	477
Detection and Auto-Configuration for Dell EqualLogic Arrays.....	477
Configuring Detection and Ports for Dell Compellent Arrays.....	478
iSCSI Optimization: Operation.....	478

Default iSCSI Optimization Values.....	479
Displaying iSCSI Optimization Information.....	479
27 Intermediate System to Intermediate System.....	481
IS-IS Protocol Overview.....	481
IS-IS Addressing.....	481
Multi-Topology IS-IS.....	482
Transition Mode.....	483
Interface Support.....	483
Adjacencies.....	483
Graceful Restart.....	483
Timers.....	484
Implementation Information.....	484
Configuration Information.....	485
Configuration Tasks for IS-IS.....	485
IS-IS Metric Styles.....	502
Configure Metric Values.....	502
Maximum Values in the Routing Table.....	502
Change the IS-IS Metric Style in One Level Only.....	502
Leaks from One Level to Another.....	504
Sample Configurations.....	505
28 Link Aggregation Control Protocol (LACP).....	510
Introduction to Dynamic LAGs and LACP.....	510
Important Points to Remember.....	510
LACP Modes.....	511
Configuring LACP Commands.....	511
LACP Configuration Tasks.....	512
Creating a LAG.....	512
Configuring the LAG Interfaces as Dynamic.....	513
Setting the LACP Long Timeout.....	513
Shared LAG State Tracking.....	514
Configuring Shared LAG State Tracking.....	515
Important Points about Shared LAG State Tracking.....	517
LACP Basic Configuration Example.....	517
Configure a LAG on ALPHA.....	517
29 Layer 2.....	526
Manage the MAC Address Table.....	526
Clearing the MAC Address Table.....	526
Setting the Aging Time for Dynamic Entries.....	526
Configuring a Static MAC Address.....	527

Displaying the MAC Address Table.....	527
MAC Learning Limit.....	527
Setting the MAC Learning Limit.....	528
mac learning-limit Dynamic.....	528
mac learning-limit station-move.....	529
Learning Limit Violation Actions.....	529
Setting Station Move Violation Actions.....	529
Recovering from Learning Limit and Station Move Violations.....	530
NIC Teaming.....	530
MAC Move Optimization.....	532
30 Link Layer Discovery Protocol (LLDP).....	533
802.1AB (LLDP) Overview.....	533
Protocol Data Units.....	533
Optional TLVs.....	534
Management TLVs.....	535
TIA-1057 (LLDP-MED) Overview.....	537
TIA Organizationally Specific TLVs.....	537
Extended Power via MDI TLV.....	541
Configure LLDP.....	542
Related Configuration Tasks.....	542
Important Points to Remember.....	542
LLDP Compatibility.....	542
CONFIGURATION versus INTERFACE Configurations.....	542
Enabling LLDP.....	543
Disabling and Undoing LLDP.....	543
Advertising TLVs.....	544
Viewing the LLDP Configuration.....	545
Viewing Information Advertised by Adjacent LLDP Agents.....	546
Configuring LLDPDU Intervals.....	547
Configuring Transmit and Receive Mode.....	548
Configuring a Time to Live.....	549
Debugging LLDP.....	549
Relevant Management Objects.....	550
31 Microsoft Network Load Balancing.....	557
NLB Unicast Mode Scenario.....	557
NLB Multicast Mode Scenario.....	558
Limitations With Enabling NLB on Switches.....	558
Benefits and Working of Microsoft Clustering.....	558
Enable and Disable VLAN Flooding	559
Configuring a Switch for NLB	559

.....	559
32 Multicast Source Discovery Protocol (MSDP).....	560
Protocol Overview.....	560
Anycast RP.....	562
Implementation Information.....	562
Configure the Multicast Source Discovery Protocol.....	563
Related Configuration Tasks.....	563
Enabling MSDP.....	567
Manage the Source-Active Cache.....	568
Viewing the Source-Active Cache.....	568
Limiting the Source-Active Cache.....	569
Clearing the Source-Active Cache.....	569
Enabling the Rejected Source-Active Cache.....	569
Accept Source-Active Messages that Fail the RFP Check.....	569
Specifying Source-Active Messages.....	573
Limiting the Source-Active Messages from a Peer.....	574
Preventing MSDP from Caching a Local Source.....	574
Preventing MSDP from Caching a Remote Source.....	575
Preventing MSDP from Advertising a Local Source.....	576
Logging Changes in Peership States.....	577
Terminating a Peership.....	577
Clearing Peer Statistics.....	578
Debugging MSDP.....	578
MSDP with Anycast RP.....	579
Configuring Anycast RP.....	580
Reducing Source-Active Message Flooding.....	581
Specifying the RP Address Used in SA Messages.....	581
MSDP Sample Configurations.....	584
33 Multiple Spanning Tree Protocol (MSTP).....	587
Protocol Overview.....	587
Spanning Tree Variations.....	588
Implementation Information.....	588
Configure Multiple Spanning Tree Protocol.....	588
Related Configuration Tasks.....	588
Enable Multiple Spanning Tree Globally.....	589
Creating Multiple Spanning Tree Instances.....	589
Influencing MSTP Root Selection.....	590
Interoperate with Non-Dell Networking OS Bridges.....	591
Changing the Region Name or Revision.....	591
Modifying Global Parameters.....	592

Enable BPDU Filtering Globally.....	593
Modifying the Interface Parameters.....	594
Configuring an EdgePort.....	595
Flush MAC Addresses after a Topology Change.....	596
MSTP Sample Configurations.....	596
Router 1 Running-ConfigurationRouter 2 Running-ConfigurationRouter 3 Running- ConfigurationSFTOS Example Running-Configuration.....	597
Debugging and Verifying MSTP Configurations.....	600
34 Multicast Features.....	603
Enabling IP Multicast.....	603
Implementation Information.....	603
First Packet Forwarding for Lossless Multicast.....	604
Multicast Policies.....	604
IPv4 Multicast Policies.....	604
Limiting the Number of Multicast Routes.....	605
Preventing a Host from Joining a Group.....	605
Rate Limiting IGMP Join Requests.....	609
Preventing a PIM Router from Forming an Adjacency.....	609
Preventing a Source from Registering with the RP.....	609
Preventing a PIM Router from Processing a Join.....	612
35 Open Shortest Path First (OSPFv2 and OSPFv3).....	613
Protocol Overview.....	613
Autonomous System (AS) Areas.....	613
Area Types.....	614
Networks and Neighbors.....	615
Router Types.....	615
Link-State Advertisements (LSAs).....	617
Router Priority and Cost.....	619
OSPF with the Dell Networking OS.....	619
Graceful Restart.....	620
Fast Convergence (OSPFv2, IPv4 Only).....	621
Multi-Process OSPFv2 (IPv4 only).....	621
RFC-2328 Compliant OSPF Flooding.....	622
OSPF ACK Packing.....	623
Setting OSPF Adjacency with Cisco Routers.....	623
Configuration Information.....	624
Configuration Task List for OSPFv2 (OSPF for IPv4).....	624
Troubleshooting OSPFv2.....	638
Configuration Task List for OSPFv3 (OSPF for IPv6).....	642
Enabling IPv6 Unicast Routing.....	643

Assigning IPv6 Addresses on an Interface.....	643
Assigning Area ID on an Interface.....	643
Assigning OSPFv3 Process ID and Router ID Globally.....	644
Configuring Stub Areas.....	644
Configuring Passive-Interface.....	644
Redistributing Routes.....	645
Configuring a Default Route.....	645
Enabling OSPFv3 Graceful Restart.....	646
Displaying Graceful Restart.....	647
OSPFv3 Authentication Using IPsec.....	648
36 Policy-based Routing (PBR).....	658
Overview.....	658
Implementing Policy-based Routing with Dell Networking OS.....	660
Configuration Task List for Policy-based Routing.....	660
PBR Exceptions (Permit).....	663
Sample Configuration.....	666
Create the Redirect-List GOLDAssign Redirect-List GOLD to Interface 2/11View	
Redirect-List GOLD.....	666
37 PIM Sparse-Mode (PIM-SM).....	671
Implementation Information.....	671
Protocol Overview.....	671
Requesting Multicast Traffic.....	671
Refuse Multicast Traffic.....	672
Send Multicast Traffic.....	672
Configuring PIM-SM.....	673
Related Configuration Tasks.....	673
Enable PIM-SM.....	673
Configuring S,G Expiry Timers.....	675
Configuring a Static Rendezvous Point.....	676
Overriding Bootstrap Router Updates.....	676
Configuring a Designated Router.....	677
Creating Multicast Boundaries and Domains.....	677
Enabling PIM-SM Graceful Restart.....	678
38 PIM Source-Specific Mode (PIM-SSM).....	679
Configure PIM-SMM.....	679
Related Configuration Tasks.....	679
Implementation Information.....	680
Important Points to Remember.....	680
Enabling PIM-SSM.....	680

Use PIM-SSM with IGMP Version 2 Hosts.....	681
Configuring PIM-SSM with IGMPv2.....	681
39 Port Monitoring.....	683
Important Points to Remember.....	683
Configuring Port Monitoring.....	684
Enabling Flow-Based Monitoring.....	686
Remote Port Mirroring.....	687
Remote Port Mirroring Example.....	687
Configuring Remote Port Mirroring.....	688
Displaying Remote-Port Mirroring Configurations.....	690
Configuring the Sample Remote Port Mirroring.....	690
Configuring the Encapsulated Remote Port Mirroring.....	693
Configuration steps for ERPM	693
ERPM Behavior on a typical Dell Networking OS	695
Decapsulation of ERPM packets at the Destination IP/ Analyzer.....	695
40 Private VLANs (PVLAN).....	697
Private VLAN Concepts.....	697
Using the Private VLAN Commands.....	698
Configuration Task List.....	699
Private VLAN Configuration Example.....	703
41 Per-VLAN Spanning Tree Plus (PVST+).....	707
Protocol Overview.....	707
Implementation Information.....	708
Configure Per-VLAN Spanning Tree Plus.....	708
Related Configuration Tasks.....	708
Enabling PVST+.....	708
Disabling PVST+.....	709
Influencing PVST+ Root Selection.....	709
Modifying Global PVST+ Parameters.....	711
Modifying Interface PVST+ Parameters.....	712
Configuring an EdgePort.....	713
PVST+ in Multi-Vendor Networks.....	714
Enabling PVST+ Extend System ID.....	714
PVST+ Sample Configurations.....	715
Enable BPDU Filtering globally.....	717
42 Quality of Service (QoS).....	719
Implementation Information.....	721
Port-Based QoS Configurations.....	721

Setting dot1p Priorities for Incoming Traffic.....	721
Honoring dot1p Priorities on Ingress Traffic.....	722
Configuring Port-Based Rate Policing.....	723
Configuring Port-Based Rate Shaping.....	723
Guidelines for Configuring ECN for Classifying and Color-Marking Packets.....	724
Sample configuration to mark non-ecn packets as “yellow” with Multiple traffic class.....	724
Classifying Incoming Packets Using ECN and Color-Marking.....	725
Sample configuration to mark non-ecn packets as “yellow” with single traffic class.....	727
Policy-Based QoS Configurations.....	729
DSCP Color Maps.....	729
Classify Traffic.....	731
Create a QoS Policy.....	736
Create Policy Maps.....	739
Enabling QoS Rate Adjustment.....	744
Enabling Strict-Priority Queueing.....	745
Weighted Random Early Detection.....	745
Creating WRED Profiles.....	746
Applying a WRED Profile to Traffic.....	747
Displaying Default and Configured WRED Profiles.....	747
Displaying WRED Drop Statistics.....	747
Displaying egress-queue Statistics.....	748
Classifying Layer 2 Traffic on Layer 3 Interfaces	748
Classifying Packets Based on a Combination of DSCP Code Points and VLAN IDs.....	749
43 Routing Information Protocol (RIP).....	751
Protocol Overview.....	751
RIPv1.....	751
RIPv2.....	752
Implementation Information.....	752
Configuration Information.....	752
Configuration Task List.....	752
RIP Configuration Example.....	759
44 Remote Monitoring (RMON).....	765
Implementation Information.....	765
Fault Recovery.....	766
Setting the rmon Alarm.....	766
Configuring an RMON Event.....	767
Configuring RMON Collection Statistics.....	768
Configuring the RMON Collection History.....	768
Enabling an RMON MIB Collection History Group.....	769

45 Rapid Spanning Tree Protocol (RSTP)	770
Protocol Overview.....	770
Configuring Rapid Spanning Tree.....	770
Related Configuration Tasks.....	770
Important Points to Remember.....	771
Configuring Interfaces for Layer 2 Mode.....	771
Enabling Rapid Spanning Tree Protocol Globally.....	771
Adding and Removing Interfaces.....	774
Modifying Global Parameters.....	775
Enable BPDU Filtering Globally.....	776
Modifying Interface Parameters.....	777
Configuring an EdgePort.....	778
Influencing RSTP Root Selection.....	779
SNMP Traps for Root Elections and Topology Changes.....	779
Configuring Fast Hellos for Link State Detection.....	779
 46 Security	 781
AAA Accounting.....	781
Configuration Task List for AAA Accounting.....	781
AAA Authentication.....	783
Configuration Task List for AAA Authentication.....	784
AAA Authorization.....	786
Privilege Levels Overview.....	786
Configuration Task List for Privilege Levels.....	787
RADIUS.....	791
RADIUS Authentication and Authorization.....	792
Configuration Task List for RADIUS.....	793
TACACS+.....	796
Configuration Task List for TACACS+.....	796
Choosing TACACS+ as the Authentication Method.....	796
Monitoring TACACS+.....	798
TACACS+ Remote Authentication and Authorization.....	798
Specifying a TACACS+ Server Host.....	799
Command Authorization.....	799
Protection from TCP Tiny and Overlapping Fragment Attacks.....	800
Enabling SCP and SSH.....	800
Using SCP with SSH to Copy a Software Image.....	801
Removing the RSA Host Keys and Zeroizing Storage	802
Configuring When to Re-generate an SSH Key	802
Configuring the SSH Server Key Exchange Algorithm.....	803
Configuring the HMAC Algorithm for the SSH Server.....	803

Configuring the SSH Server Cipher List.....	804
Secure Shell Authentication.....	805
Troubleshooting SSH.....	807
Telnet.....	808
VTY Line and Access-Class Configuration.....	808
VTY Line Local Authentication and Authorization.....	808
VTY Line Remote Authentication and Authorization.....	809
VTY MAC-SA Filter Support.....	810
Role-Based Access Control.....	810
Overview of RBAC.....	811
User Roles.....	814
AAA Authentication and Authorization for Roles.....	817
Role Accounting.....	820
Display Information About User Roles.....	822
47 Service Provider Bridging.....	824
VLAN Stacking.....	824
Important Points to Remember.....	825
Configure VLAN Stacking.....	826
Creating Access and Trunk Ports.....	826
Enable VLAN-Stacking for a VLAN.....	827
Configuring the Protocol Type Value for the Outer VLAN Tag.....	827
Configuring Options for Trunk Ports.....	828
Debugging VLAN Stacking.....	829
VLAN Stacking in Multi-Vendor Networks.....	829
VLAN Stacking Packet Drop Precedence.....	833
Enabling Drop Eligibility.....	833
Honoring the Incoming DEI Value.....	834
Marking Egress Packets with a DEI Value.....	834
Dynamic Mode CoS for VLAN Stacking.....	835
Mapping C-Tag to S-Tag dot1p Values.....	836
Layer 2 Protocol Tunneling.....	837
Implementation Information.....	839
Enabling Layer 2 Protocol Tunneling.....	840
Specifying a Destination MAC Address for BPDUs.....	840
Setting Rate-Limit BPDUs.....	840
Debugging Layer 2 Protocol Tunneling.....	841
Provider Backbone Bridging.....	841
48 sFlow.....	842
Overview.....	842
Implementation Information.....	842

Important Points to Remember.....	843
Enabling and Disabling sFlow.....	843
Enabling and Disabling sFlow on an Interface.....	843
Enabling sFlow Max-Header Size Extended.....	844
sFlow Show Commands.....	845
Displaying Show sFlow Global.....	845
Displaying Show sFlow on an Interface.....	846
Displaying Show sFlow on a Stack Unit.....	846
Configuring Specify Collectors.....	846
Changing the Polling Intervals.....	847
Changing the Sampling Rate.....	847
Sub-Sampling.....	848
Back-Off Mechanism.....	848
sFlow on LAG ports.....	849
Enabling Extended sFlow.....	849
49 Simple Network Management Protocol (SNMP).....	851
Implementation Information.....	851
Configuration Task List for SNMP.....	851
Important Points to Remember.....	852
SNMPv3 Compliance With FIPS.....	852
Set up SNMP.....	853
Creating a Community.....	854
Setting Up User-Based Security (SNMPv3).....	854
Reading Managed Object Values.....	856
Writing Managed Object Values.....	857
Configuring Contact and Location Information using SNMP.....	857
Subscribing to Managed Object Value Updates using SNMP.....	858
Enabling a Subset of SNMP Traps.....	859
Enabling an SNMP Agent to Notify Syslog Server Failure.....	861
Copy Configuration Files Using SNMP.....	862
Copying a Configuration File.....	864
Copying Configuration Files via SNMP.....	865
Copying the Startup-Config Files to the Running-Config.....	866
Copying the Startup-Config Files to the Server via FTP.....	866
Copying the Startup-Config Files to the Server via TFTP.....	867
Copying a Binary File to the Startup-Configuration.....	867
Additional MIB Objects to View Copy Statistics.....	868
MIB Support to Display the Available Memory Size on Flash.....	869
Viewing the Available Flash Memory Size.....	869
MIB Support to Display the Software Core Files Generated by the System.....	869
Viewing the Software Core Files Generated by the System.....	870

Obtaining a Value for MIB Objects.....	870
Manage VLANs using SNMP.....	871
Creating a VLAN.....	871
Assigning a VLAN Alias.....	871
Displaying the Ports in a VLAN.....	872
Add Tagged and Untagged Ports to a VLAN.....	873
Enabling and Disabling a Port using SNMP.....	874
Fetch Dynamic MAC Entries using SNMP.....	875
Deriving Interface Indices.....	877
Monitor Port-Channels.....	878
BMP Functionality Using SNMP SET.....	879
Entity MIBS.....	880
Physical Entity.....	880
Containment Tree.....	880
Troubleshooting SNMP Operation.....	881
50 Stacking.....	882
Stacking MXL 10/40GbE Switches.....	882
Stack Management Roles.....	883
Stack Master Election.....	884
Failover Roles.....	885
MAC Addressing.....	885
Stacking LAG.....	885
Supported Stacking Topologies.....	885
Stack Group/Port Numbers.....	887
Configuring a Switch Stack.....	888
Stacking Prerequisites.....	888
Master Selection Criteria.....	889
Configuring Priority and stack-group.....	889
Cabling Stacked Switches.....	890
Accessing the CLI.....	891
Configuring and Bringing Up a Stack.....	891
Removing a Switch from a Stack.....	894
Adding a Stack Unit.....	895
Merging Two Stacks.....	896
Splitting a Stack.....	896
Managing Redundant Stack Management.....	896
Resetting a Unit on a Stack.....	897
Verify a Stack Configuration.....	897
Using Show Commands.....	898
Troubleshooting a Switch Stack.....	900
Failure Scenarios.....	902

Stack Member FailsUnplugged Stacking CableMaster Switch FailsStack-Link Flapping ErrorMaster Switch Recovers from FailureStack Unit in Card-Problem State Due to Incorrect Dell Networking OS VersionStack Unit in Card-Problem State Due to Configuration Mismatch.....	902
Upgrading a Switch Stack.....	905
Upgrading a Single Stack Unit.....	906
51 Storm Control.....	908
Configure Storm Control.....	908
Configuring Storm Control from INTERFACE Mode.....	908
Configuring Storm Control from CONFIGURATION Mode.....	908
52 Spanning Tree Protocol (STP).....	909
Protocol Overview.....	909
Configure Spanning Tree.....	909
Related Configuration Tasks.....	909
Important Points to Remember.....	910
Configuring Interfaces for Layer 2 Mode.....	911
Enabling Spanning Tree Protocol Globally.....	912
Adding an Interface to the Spanning Tree Group.....	915
Removing an Interface from the Spanning Tree Group.....	915
Modifying Global Parameters.....	915
Modifying Interface STP Parameters.....	916
Enabling PortFast.....	917
Prevent Network Disruptions with BPDU Guard.....	918
Global BPDU Filtering.....	920
Interface BPDU Filtering.....	920
Selecting STP Root.....	921
STP Root Guard.....	922
Root Guard Scenario.....	922
Configuring Root Guard.....	923
SNMP Traps for Root Elections and Topology Changes.....	924
Displaying STP Guard Configuration.....	924
53 System Time and Date.....	925
Network Time Protocol.....	925
Protocol Overview.....	926
Configure the Network Time Protocol.....	927
Enabling NTP.....	927
Configuring NTP Broadcasts.....	928
Disabling NTP on an Interface.....	928
Configuring a Source IP Address for NTP Packets.....	928

Configuring NTP Authentication.....	929
Dell Networking OS Time and Date.....	931
Configuration Task List	931
Set Daylight Saving Time.....	933
54 Tunneling	936
Configuring a Tunnel.....	936
Configuring Tunnel keepalive.....	937
Configuring the ip and ipv6 unnumbered.....	938
Configuring the Tunnel allow-remote.....	938
Configuring the Tunnel Source Anylocal.....	939
55 Uplink Failure Detection (UFD).....	940
Feature Description.....	940
How Uplink Failure Detection Works.....	941
UFD and NIC Teaming.....	942
Important Points to Remember.....	943
Configuring Uplink Failure Detection.....	943
Clearing a UFD-Disabled Interface.....	945
Displaying Uplink Failure Detection.....	946
Sample Configuration: Uplink Failure Detection.....	948
56 Upgrade Procedures.....	950
Get Help with Upgrades.....	950
57 Virtual LANs (VLANs).....	951
Default VLAN.....	951
Port-Based VLANs.....	952
VLANs and Port Tagging.....	953
Configuration Task List.....	953
Configuring Native VLANs.....	957
Enabling Null VLAN as the Default VLAN.....	958
58 Virtual Link Trunking (VLT).....	959
Overview.....	959
Multi-domain VLT.....	960
VLT Terminology.....	961
Configure Virtual Link Trunking.....	961
Important Points to Remember.....	962
Configuration Notes.....	962
RSTP and VLT.....	966
VLT Bandwidth Monitoring.....	967

VLT and IGMP Snooping.....	967
VLT Port Delayed Restoration.....	967
PIM-Sparse Mode Support on VLT.....	968
VLT Multicast.....	969
VLT Unicast Routing.....	970
Non-VLT ARP Sync.....	971
RSTP Configuration.....	972
Preventing Forwarding Loops in a VLT Domain.....	972
Sample RSTP Configuration.....	972
Configuring VLT.....	973
Configuring a VLT Interconnect.....	973
Configuring a VLT Backup Link.....	974
Configuring a VLT Port Delay Period.....	975
Reconfiguring the Default VLT Settings (Optional)	975
Connecting a VLT Domain to an Attached Access Device (Switch or Server).....	976
Configuring a VLT VLAN Peer-Down (Optional).....	977
Configure Multi-domain VLT (mVLT) (Optional).....	977
Verifying a VLT Configuration.....	980
Connecting a VLT Domain.....	983
PVST+ Configuration.....	988
Sample PVST+ Configuration.....	988
mVLT Configuration Example.....	989
In Domain 1, configure the VLT domain and VLTi on Peer 1Configure mVLT on Peer 1Add links to the mVLT port-channel on Peer 1Next, configure the VLT domain and VLTi on Peer 2Configure mVLT on Peer 2Add links to the mVLT port-channel on Peer 2In Domain 2, configure the VLT domain and VLTi on Peer 3Configure mVLT on Peer 3Add links to the mVLT port-channel on Peer 3Configure the VLT domain and VLTi on Peer 4Configure mVLT on Peer 4Add links to the mVLT port-channel on Peer 4.....	990
PIM-Sparse Mode Configuration Example.....	991
Additional VLT Sample Configurations.....	992
Configuring Virtual Link Trunking (VLT Peer 1)Configuring Virtual Link Trunking (VLT Peer 2)Verifying a Port-Channel Connection to a VLT Domain (From an Attached Access Switch).....	992
Troubleshooting VLT.....	994
Specifying VLT Nodes in a PVLAN.....	996
Association of VLTi as a Member of a PVLAN.....	997
MAC Synchronization for VLT Nodes in a PVLAN.....	998
PVLAN Operations When One VLT Peer is Down.....	998
PVLAN Operations When a VLT Peer is Restarted.....	998
Interoperation of VLT Nodes in a PVLAN with ARP Requests.....	998
Scenarios for VLAN Membership and MAC Synchronization With VLT Nodes in PVLAN.....	999
Configuring a VLT VLAN or LAG in a PVLAN.....	1000

Creating a VLT LAG or a VLT VLAN.....	1001
Associating the VLT LAG or VLT VLAN in a PVLAN.....	1002
Proxy ARP Capability on VLT Peer Nodes.....	1003
Working of Proxy ARP for VLT Peer Nodes.....	1003
Configuring VLAN-Stack over VLT.....	1004
59 Virtual Router Redundancy Protocol (VRRP).....	1008
VRRP Overview.....	1008
VRRP Benefits.....	1009
VRRP Implementation.....	1009
VRRP Configuration.....	1010
Configuration Task List.....	1010
Setting VRRP Initialization Delay.....	1019
Sample Configurations.....	1020
VRRP for an IPv4 Configuration.....	1020
60 Standards Compliance.....	1023
IEEE Compliance.....	1023
RFC and I-D Compliance.....	1024
General Internet Protocols.....	1024
General IPv4 Protocols.....	1025
Border Gateway Protocol (BGP).....	1026
Open Shortest Path First (OSPF).....	1026
Routing Information Protocol (RIP).....	1027
Network Management.....	1027
MIB Location.....	1030
61 FC Flex IO Modules.....	1031
FC Flex IO Modules.....	1031
Understanding and Working of the FC Flex IO Modules.....	1031
FC Flex IO Modules Overview.....	1031
FC Flex IO Module Capabilities and Operations.....	1033
Guidelines for Working with FC Flex IO Modules.....	1033
Processing of Data Traffic.....	1036
Installing and Configuring the Switch.....	1037
Interconnectivity of FC Flex IO Modules with Cisco MDS Switches.....	1040
Data Center Bridging (DCB).....	1041
Ethernet Enhancements in Data Center Bridging.....	1042
Enabling Data Center Bridging.....	1049
QoS dot1p Traffic Classification and Queue Assignment.....	1050
Configure Enhanced Transmission Selection.....	1051
Configure a DCBx Operation.....	1053

Verifying the DCB Configuration.....	1063
PFC and ETS Configuration Examples.....	1073
Using PFC and ETS to Manage Data Center Traffic.....	1073
Fibre Channel over Ethernet for FC Flex IO Modules.....	1077
NPIV Proxy Gateway for FC Flex IO Modules.....	1077
NPIV Proxy Gateway Configuration on FC Flex IO Modules	1077
NPIV Proxy Gateway Operations and Capabilities.....	1078
Configuring an NPIV Proxy Gateway.....	1082
Displaying NPIV Proxy Gateway Information.....	1088

About this Guide

This guide describes the supported protocols and software features, and provides configuration instructions and examples, for the Dell Networking MXL 10/40GbE Switch IO Module.

The MXL 10/40GbE Switch IO Module is installed in a Dell PowerEdge M1000e Enclosure. For information about how to install and perform the initial switch configuration, refer to the Getting Started Guides on the Dell Support website at <http://support.dell.com/manuals>.

Though this guide contains information on protocols, it is not intended to be a complete reference. This guide is a reference for configuring protocols on Dell Networking systems. For complete information about protocols, refer to related documentation, including IETF requests for comments (RFCs). The instructions in this guide cite relevant RFCs. The [Standards Compliance](#) chapter contains a complete list of the supported RFCs and management information base files (MIBs).

Audience

This document is intended for system administrators who are responsible for configuring and maintaining networks and assumes knowledge in Layer 2 and Layer 3 networking technologies.

Conventions

This guide uses the following conventions to describe command syntax.

Keyword	Keywords are in Courier (a monospaced font) and must be entered in the CLI as listed.
<i>parameter</i>	Parameters are in italics and require a number or word to be entered in the CLI.
{X}	Keywords and parameters within braces must be entered in the CLI.
[X]	Keywords and parameters within brackets are optional.
x y	Keywords and parameters separated by a bar require you to choose one option.
x y	Keywords and parameters separated by a double bar allows you to choose any or all of the options.

Information Symbols

This book uses the following information symbols.



NOTE: The Note icon signals important operational information.



CAUTION: The Caution icon signals information about situations that could result in equipment damage or loss of data.



WARNING: The Warning icon signals information about hardware handling that could result in injury.

* (Exception). This symbol is a note associated with additional text on the page that is marked with an asterisk.

Related Documents

For more information about the Dell Networking MXL 10/40GbE Switch IO Module, refer to the following documents:


- *Dell Networking OS Command Reference*
- *Dell Quick Start Guide*
- *Dell Networking OS Release Notes*

Configuration Fundamentals

The Dell Networking operating system command line interface (CLI) is a text-based interface you can use to configure interfaces and protocols.

The CLI is structured in modes for security and management purposes. Different sets of commands are available in each mode, and you can limit user access to modes using privilege levels.

In the Dell Networking OS, after you enable a command, it is entered into the running configuration file. You can view the current configuration for the whole system or for a particular CLI mode. To save the current configuration, copy the running configuration to another location. For more information, refer to [Save the Running-Configuration](#).

 **NOTE:** You can use the chassis management controller (CMC) out-of-band management interface to access and manage an MXL Switch using the CLI. For information about how to access the CMC to configure an MXL Switch, refer to the *Dell Chassis Management Controller (CMC) User's Guide* on the Dell Support website.

Accessing the Command Line

Access the CLI through a serial console port or a Telnet session.

When the system successfully boots, enter the command line in EXEC mode.

```
telnet 172.31.1.53
Trying 172.31.1.53...
Connected to 172.31.1.53.
Escape character is '^]'.
Login: username
Password:
Dell>
```

CLI Modes

Different sets of commands are available in each mode.

A command found in one mode cannot be executed from another mode (except for EXEC mode commands with a preceding `do` command (refer to [The do Command](#) section)).


You can set user access rights to commands and command modes using privilege levels; for more information about privilege levels and security options, refer to the *Privilege Levels Overview* section in the [Security](#) chapter.

The CLI is divided into three major mode levels:

- **EXEC mode** is the default mode and has a privilege level of 1, which is the most restricted level. Only a limited selection of commands is available, notably the `show` commands, which allow you to view system information.
- **EXEC Privilege mode** has commands to view configurations, clear counters, manage configuration files, run diagnostics, and enable or disable debug operations. The privilege level is 15, which is unrestricted. You can configure a password for this mode; refer to the *Configure the Enable Password* section in the [Getting Started](#) chapter.
- **CONFIGURATION mode** allows you to configure security features, time settings, set logging and SNMP functions, configure static ARP and MAC addresses, and set line cards on the system.

Beneath CONFIGURATION mode are submodes that apply to interfaces, protocols, and features. The following example shows the submode command structure. Two sub-CONFIGURATION modes are important when configuring the chassis for the first time:

- **INTERFACE sub-mode** is the mode in which you configure Layer 2 and Layer 3 protocols and IP services specific to an interface. An interface can be physical (Management interface, 10 Gigabit Ethernet, 40 Gigabit Ethernet, or synchronous optical network technologies [SONET]) or logical (Loopback, Null, port channel, or virtual local area network [VLAN]).
- **LINE sub-mode** is the mode in which you to configure the console and virtual terminal lines.

 **NOTE:** At any time, entering a question mark (?) displays the available command options. For example, when you are in CONFIGURATION mode, entering the question mark first lists all available commands, including the possible submodes.

The CLI modes are:

```
EXEC
EXEC Privilege
CONFIGURATION
  INTERFACE
    TEN GIGABIT ETHERNET
    FORTY GIGABIT ETHERNET
    INTERFACE RANGE
    LOOPBACK
    MANAGEMENT ETHERNET
    MONITOR SESSION
    NULL
    PORT-CHANNEL
    VLAN
  IP
  IP ACCESS-LIST
    STANDARD ACCESS-LIST
    EXTENDED ACCESS-LIST
  LINE
    CONSOLE
    VIRTUAL TERMINAL
  MAC ACCESS-LIST
  MONITOR SESSION
  MULTIPLE SPANNING TREE
  PROTOCOL GVRP
  PROTOCOL LLDP
  PER-VLAN SPANNING TREE
  RAPID SPANNING TREE
  ROUTE-MAP
  ROUTER OSPF
```


Navigating CLI Modes

The Dell Networking OS prompt changes to indicate the CLI mode.

The following table lists the CLI mode, its prompt, and information about how to access and exit the CLI mode. Move linearly through the command modes, except for the `end` command which takes you directly to EXEC Privilege mode and the `exit` command which moves you up one command mode level.



 **NOTE:** Sub-CONFIGURATION modes all have the letters “conf” in the prompt with more modifiers to identify the mode and slot/port information.

Table 1. Dell Networking OS Command Modes

CLI Command Mode	Prompt	Access Command
EXEC	Dell>	Access the router through the console or Telnet.
EXEC Privilege	Dell#	<ul style="list-style-type: none"> From EXEC mode, enter the <code>enable</code> command. From any other mode, use the <code>end</code> command.
CONFIGURATION	Dell (conf) #	<ul style="list-style-type: none"> From EXEC privilege mode, enter the <code>configure</code> command. From every mode except EXEC and EXEC Privilege, enter the <code>exit</code> command.

 **NOTE:** Access all of the following modes from CONFIGURATION mode.

AS-PATH ACL	Dell (config-as-path) #	<code>ip as-path access-list</code>
Gigabit Ethernet Interface	Dell (conf-if-gi-0/0) #	<code>interface (INTERFACE modes)</code>
10 Gigabit Ethernet Interface	Dell (conf-if-te-0/0) #	<code>interface (INTERFACE modes)</code>
Interface Range	Dell (conf-if-range) #	<code>interface (INTERFACE modes)</code>
Loopback Interface	Dell (conf-if-lo-0) #	<code>interface (INTERFACE modes)</code>
Management Ethernet Interface	Dell (conf-if-ma-0/0) #	<code>interface (INTERFACE modes)</code>
Null Interface	Dell (conf-if-nu-0) #	<code>interface (INTERFACE modes)</code>
Port-channel Interface	Dell (conf-if-po-0) #	<code>interface (INTERFACE modes)</code>
Tunnel Interface	Dell (conf-if-tu-0) #	<code>interface (INTERFACE modes)</code>
VLAN Interface	Dell (conf-if-vl-0) #	<code>interface (INTERFACE modes)</code>

CLI Command Mode	Prompt	Access Command
STANDARD ACCESS-LIST	Dell(config-std-nacl)#	ip access-list standard (IP ACCESS-LIST Modes)
EXTENDED ACCESS-LIST	Dell(config-ext-nacl)#	ip access-list extended (IP ACCESS-LIST Modes)
IP COMMUNITY-LIST	Dell(config-community-list)#	ip community-list
AUXILIARY	Dell(config-line-aux)#	line (LINE Modes)
CONSOLE	Dell(config-line-console)#	line (LINE Modes)
VIRTUAL TERMINAL	Dell(config-line-vty)#	line (LINE Modes)
STANDARD ACCESS-LIST	Dell(config-std-macl)#	mac access-list standard (MAC ACCESS-LIST Modes)
EXTENDED ACCESS-LIST	Dell(config-ext-macl)#	mac access-list extended (MAC ACCESS-LIST Modes)
MULTIPLE SPANNING TREE	Dell(config-mstp)#	protocol spanning-tree mstp
Per-VLAN SPANNING TREE Plus	Dell(config-pvst)#	protocol spanning-tree pvst
PREFIX-LIST	Dell(conf-nprefixl)#	ip prefix-list
RAPID SPANNING TREE	Dell(config-rstp)#	protocol spanning-tree rstp
REDIRECT	Dell(conf-redirect-list)#	ip redirect-list
ROUTE-MAP	Dell(config-route-map)#	route-map
ROUTER BGP	Dell(conf-router_bgp)#	router bgp
BGP ADDRESS-FAMILY	Dell(conf-router_bgp_af)# (for IPv4) Dell(conf-routerZ_bgpv6_af)# (for IPv6)	address-family {ipv4 multicast ipv6 unicast} (ROUTER BGP Mode)
ROUTER ISIS	Dell(conf-router_isis)#	router isis
ISIS ADDRESS-FAMILY	Dell(conf-router_isis-af_ipv6)#	address-family ipv6 unicast (ROUTER ISIS Mode)
ROUTER OSPF	Dell(conf-router_ospf)#	router ospf
ROUTER OSPFV3	Dell(conf-ipv6router_ospf)#	ipv6 router ospf
ROUTER RIP	Dell(conf-router_rip)#	router rip
SPANNING TREE	Dell(config-span)#	protocol spanning-tree 0

CLI Command Mode	Prompt	Access Command
TRACE-LIST	Dell (conf-trace-acl) #	ip trace-list
CLASS-MAP	Dell (config-class-map) #	class-map
CONTROL-PLANE	Dell (conf-control-cpuqos) #	control-plane-cpuqos
DCB POLICY	Dell (conf-dcb-in) # (for input policy) Dell (conf-dcb-out) # (for output policy)	dcb-input for input policy dcb-output for output policy
DHCP	Dell (config-dhcp) #	ip dhcp server
DHCP POOL	Dell (config-dhcp-pool-name) #	pool (DHCP Mode)
ECMP	Dell (conf-ecmp-group-ecmp-group-id) #	ecmp-group
EIS	Dell (conf-mgmt-eis) #	management egress-interface-selection
FRRP	Dell (conf-frrp-ring-id) #	protocol frrp
LLDP	Dell (conf-lldp) # or Dell (conf-if-interface-lldp) #	protocol lldp (CONFIGURATION or INTERFACE Modes)
LLDP MANAGEMENT INTERFACE	Dell (conf-lldp-mgmtIf) #	management-interface (LLDP Mode)
LINE	Dell (config-line-console) or Dell (config-line-vty)	line console or line vty
MONITOR SESSION	Dell (conf-mon-sess-sessionID) #	monitor session
OPENFLOW INSTANCE	Dell (conf-of-instance-of-id) #	openflow of-instance
PORT-CHANNEL FAILOVER-GROUP	Dell (conf-po-failover-grp) #	port-channel failover-group
PRIORITY GROUP	Dell (conf-pg) #	priority-group
PROTOCOL GVRP	Dell (config-gvrp) #	protocol gvrp
QOS POLICY	Dell (conf-qos-policy-outputs) #	qos-policy-output
VLT DOMAIN	Dell (conf-vlt-domain) #	vlt domain
VRRP	Dell (conf-if-interface-type-slot/port-vrid-vrrp-group-id) #	vrrp-group

CLI Command Mode	Prompt	Access Command
u-Boot	Dell (=>) #	Press any key when the following line appears on the console during a system boot: Hit any key to stop autoboot:
UPLINK STATE GROUP	Dell (conf-uplink-state-group- <i>groupID</i>) #	uplink-state-group

The following example shows how to change the command mode from CONFIGURATION mode to PROTOCOL SPANNING TREE.

Example of Changing Command Modes

```
Dell(conf)#protocol spanning-tree 0
Dell(config-span)#
```

The do Command

You can enter an EXEC mode command from any CONFIGURATION mode (CONFIGURATION, INTERFACE, SPANNING TREE, and so on.) without having to return to EXEC mode by preceding the EXEC mode command with the `do` command.

The following example shows the output of the `do` command: `enable`, `disable`, `exit`, and `configure`.

```
Dell(conf)#do show system brief

Stack MAC : 00:1e:c9:f1:04:22


Reload Type : normal-reload [Next boot : normal-reload]

-- Stack Info --
Unit UnitType      Status      ReqTyp      CurTyp      Version  Ports
-----
0    Management  online     MXL-10/40GbE  MXL-10/40GbE  8-3-16-47  56
1    Member      not present
2    Member      not present
3    Member      not present
4    Member      not present
5    Member      not present
```

Undoing Commands

When you enter a command, the command line is added to the running configuration file (running-config).

To disable a command and remove it from the running-config, enter the `no` command, then the original command. For example, to delete an IP address configured on an interface, use the `no ip address ip-address` command.

 **NOTE:** Use the help or `?` command as described in [Obtaining Help](#).

The first bold line shows the assigned IP address, the second bold line shows the `no` form of the IP address command, and the last bold line shows the IP address removed.

Example of Viewing Disabled Commands

```
Dell(conf)#interface gigabitethernet 4/17
Dell(conf-if-gi-4/17)#ip address 192.168.10.1/24
Dell(conf-if-gi-4/17)#show config
!
 interface GigabitEthernet 4/17
  ip address 192.168.10.1/24
no shutdown
Dell(conf-if-gi-4/17)#no ip address
Dell(conf-if-gi-4/17)#show config
!
interface GigabitEthernet 4/17
  no ip address
  no shutdown
```

Layer 2 protocols are disabled by default. To enable Layer 2 protocols, use the `no disable` command. For example, in PROTOCOL SPANNING TREE mode, enter `no disable` to enable Spanning Tree.

Obtaining Help

Obtain a list of keywords and a brief functional description of those keywords at any CLI mode using the `?` or help command:

- To list the keywords available in the current mode, enter `?` at the prompt or after a keyword.
- Enter `?` after a prompt lists all of the available keywords. The output of this command is the same for the help command.

```
Dell#?
start      Start Shell
capture    Capture Packet
cd         Change current directory
clear      Reset functions
clock      Manage the system clock
configure  Configuring from terminal
copy       Copy from one file to another
--More--
```

- Enter `?` after a partial keyword lists all of the keywords that begin with the specified letters.

```
Dell(conf)#cl?
class-map
clock
Dell(conf)#cl
```

- Enter `[space]?` after a keyword lists all of the keywords that can follow the specified keyword.

```
Dell(conf)#clock ?
summer-time    Configure summer (daylight savings) time
timezone       Configure time zone
Dell(conf)#clock
```

Entering and Editing Commands

Notes for entering commands.

- The CLI is not case-sensitive.
- You can enter partial CLI keywords.
 - Enter the minimum number of letters to uniquely identify a command. For example, you cannot enter `cl` as a partial keyword because both the `clock` and `class-map` commands begin with the letters “cl.” You can enter `clo`, however, as a partial keyword because only one command begins with those three letters.
- The TAB key auto-completes keywords in commands. Enter the minimum number of letters to uniquely identify a command.
- The UP and DOWN arrow keys display previously entered commands (refer to [Command History](#)).
- The BACKSPACE and DELETE keys erase the previous letter.
- Key combinations are available to move quickly across the command line. The following list describes these short-cut key combinations.

Short-Cut Key Combination	Action
CNTL-A	Moves the cursor to the beginning of the command line.
CNTL-B	Moves the cursor back one character.
CNTL-D	Deletes character at cursor.
CNTL-E	Moves the cursor to the end of the line.
CNTL-F	Moves the cursor forward one character.
CNTL-I	Completes a keyword.
CNTL-K	Deletes all characters from the cursor to the end of the command line.
CNTL-L	Re-enters the previous command.
CNTL-N	Return to more recent commands in the history buffer after recalling commands with CTRL-P or the UP arrow key.
CNTL-P	Recalls commands, beginning with the last command.
CNTL-R	Re-enters the previous command.
CNTL-U	Deletes the line.
CNTL-W	Deletes the previous word.
CNTL-X	Deletes the line.
CNTL-Z	Ends continuous scrolling of command outputs.
Esc B	Moves the cursor back one word.
Esc F	Moves the cursor forward one word.
Esc D	Deletes all characters from the cursor to the end of the word.

Command History

The Dell Networking OS maintains a history of previously-entered commands for each mode. For example:

- When you are in EXEC mode, the UP and DOWN arrow keys display the previously-entered EXEC mode commands.
- When you are in CONFIGURATION mode, the UP or DOWN arrows keys recall the previously-entered CONFIGURATION mode commands.

Filtering `show` Command Outputs

Filter the output of a `show` command to display specific information by adding `| [except | find | grep | no-more | save] specified_text` after the command.


The variable `specified_text` is the text for which you are filtering and it IS case sensitive unless you use the `ignore-case` sub-option.

Starting with the Dell Networking OS version 7.8.1.0, the `grep` command accepts an `ignore-case` sub-option that forces the search to case-insensitive. For example, the commands:

- `show run | grep Ethernet` returns a search result with instances containing a capitalized "Ethernet," such as `interface GigabitEthernet 0/0`.
- `show run | grep ethernet` does not return that search result because it only searches for instances containing a non-capitalized "ethernet."
- `show run | grep Ethernet ignore-case` returns instances containing both "Ethernet" and "ethernet."

The `grep` command displays only the lines containing specified text. The following shows this command used in combination with the `do show stack-unit all stack-ports pfc details | grep 0` command.

```
Dell(conf)#do show stack-unit all stack-ports all pfc details | grep 0
stack unit 0 stack-port all
 0 Pause Tx pkts, 0 Pause Rx pkts
 0 Pause Tx pkts, 0 Pause Rx pkts
 0 Pause Tx pkts, 0 Pause Rx pkts
 0 Pause Tx pkts, 0 Pause Rx pkts
 0 Pause Tx pkts, 0 Pause Rx pkts
 0 Pause Tx pkts, 0 Pause Rx pkts
```

 **NOTE:** The Dell Networking OS accepts a space or no space before and after the pipe. To filter a phrase with spaces, underscores, or ranges, enclose the phrase with double quotation marks.

The `except` keyword displays text that does not match the specified text. The following example shows this command used in combination with the `do show stack-unit all stack-ports all pfc details | except 0` command.

Example of the `except` Keyword

Example of the `find` Keyword

```
Dell(conf)#do show stack-unit all stack-ports all pfc details | except 0

Admin mode is On
Admin is enabled
Local is enabled
Link Delay 45556 pause quantum

stack unit 1 stack-port all

Admin mode is On
Admin is enabled
```


The `find` keyword displays the output of the `show` command beginning from the first occurrence of specified text.

```
Dell(conf)#do show stack-unit all stack-ports all pfc details | find 0
stack unit 0 stack-port all
Admin mode is On
Admin is enabled
Local is enabled
Link Delay 45556 pause quantum
0 Pause Tx pkts, 0 Pause Rx pkts

stack unit 1 stack-port all
```

The `no-more` command displays the output all at once rather than one screen at a time. This is similar to the `terminal length` command except that the `no-more` option affects the output of the specified command only.

The `save` command copies the output to a file for future reference.

 **NOTE:** You can filter a single command output multiple times. The `save` option must be the last option entered. For example: `Dell# command | grep regular-expression | except regular-expression | grep other-regular-expression | find regular-expression | save.`

Multiple Users in Configuration Mode

Dell Networking OS notifies all users when there are multiple users logged in to CONFIGURATION mode.

A warning message indicates the username, type of connection (console or VTY), and in the case of a VTY connection, the IP address of the terminal on which the connection was established. For example:

- On the system that telnets into the switch, this message appears:
% Warning: The following users are currently configuring the system:
User "<username>" on line console0
- On the system that is connected over the console, this message appears:
% Warning: User "<username>" on line vty0 "10.11.130.2" is in configuration mode

If either of these messages appears, Dell Networking recommends coordinating with the users listed in the message so that you do not unintentionally overwrite each other's configuration changes.

Getting Started

This chapter describes how you start configuring your system.

When you power up the chassis, the system performs a power-on self test (POST) during which the route processor module (RPM), switch fabric module (SFM), and line card status light emitting diodes (LEDs) blink green. The system then loads the Dell Networking operating system. Boot messages scroll up the terminal window during this process. No user interaction is required if the boot process proceeds without interruption.

When the boot process completes, the RPM and line card status LEDs remain online (green) and the console monitor displays the EXEC mode prompt.

For details about using the command line interface (CLI), refer to the [Accessing the Command Line](#) section in the [Configuration Fundamentals](#) chapter.

Console Access

The MXL 10/40GbE Switch IO Module has two management ports available for system access: a serial console port and an out-of-bounds (OOB) port.

Serial Console

A universal serial bus (USB) (A-Type) connector is located at the front panel. The USB can be defined as an External Serial Console (RS-232) port, and is labeled on the MXL 10/40GbE Switch IO Module chassis. The USB is present on the lower side, as you face the I/O side of the chassis, as shown.

Serial Console



External Serial Port with a USB Connector

The following table lists the pin assignments.

Table 2. Pin Assignments

USB Pin Number	Signal Name
Pin 1	RTS
Pin 2	RX
Pin 3	TX
Pin 4	CTS
Pin 5, 6	GND
RxD	Chassis GND

Accessing the CLI Interface and Running Scripts Using SSH

In addition to the capability to access a device using a console connection or a Telnet session, you can also use SSH for secure, protected communication with the device. You can open an SSH session and run commands or script files. This method of connectivity is supported with MXL switch and provides a reliable, safe communication mechanism.

Entering CLI commands Using an SSH Connection

You can run CLI commands by entering any one of the following syntax to connect to a switch using the preconfigured user credentials using SSH:

```
ssh username@hostname <CLI Command>
```

or

```
echo <CLI Command> | ssh admin@hostname
```

The SSH server transmits the terminal commands to the CLI shell and the results are displayed on the screen non-interactively.

Executing Local CLI Scripts Using an SSH Connection

You can execute CLI commands by entering a CLI script in one of the following ways:

```
ssh username@hostname <CLIScript.file>
```

or

```
cat < CLIScript.file > | ssh admin@hostname
```

The script is run and the actions contained in the script are performed.

Following are the points to remember, when you are trying to establish an SSH session to the device to run commands or script files:

- There is an upper limit of 10 concurrent sessions in SSH. Therefore, you might expect a failure in executing SSH-related scripts.
- To avoid denial of service (DoS) attacks, a rate-limit of 10 concurrent sessions per minute in SSH is devised. Therefore, you might experience a failure in executing SSH-related scripts when multiple short SSH commands are executed.
- If you issue an interactive command in the SSH session, the behavior may not really be interactive.
- In some cases, when you use an SSH session, when certain show commands such as `show tech-support` produce large volumes of output, sometimes few characters from the output display are truncated and not displayed. This may cause one of the commands to fail for syntax error. In such cases, if you add few newline characters before the failed command, the output displays completely.

Execution of commands on CLI over SSH does not notice the errors that have occurred while executing the command. As a result, you cannot identify, whether a command has failed to be processed. The console output though is redirected back over SSH.

Boot Process

After you follow the *Installation Procedure* in the *Getting Started Guide*, the MXL switch boots up.

The MXL switch with the Dell Networking OS version 8.3.16.1 requires boot flash version 4.0.1.0 and boot selector version 4.0.0.0. The following example shows the completed boot process.

```
syncing disks... done
unmounting file systems...
unmounting /f10/flash (/dev/ld0e)...
unmounting /usr (mfs:31)...
unmounting /lib (mfs:23)...
unmounting /f10 (mfs:20)...
unmounting /tmp (mfs:15)...
unmounting /kern (kernfs)...
unmounting / (/dev/md0a)... done
rebooting...
```

```
NetLogic XLP Stage 1 Loader
Built by build at tools-sjc-01 on Thu May 31 23:53:38 2012
IOM Boot Selector Label 4.0.0.0
```

```
Nodes online: 1
  GPIO 22 init'ed as an output
  GPIO 23 init'ed as an output
I2C0 speed = 30 KHz, prescaler = 0x0377.
Initialized I2C0 Controller.
I2C1 speed = 100 KHz, prescaler = 0x0109.
Initialized I2C1 Controller.
DDR SPD: Node 0 Channel 0 Mem size = 2048 MB
DDR SPD: Node 0 DRAM frequency 666 MHz
DDR SPD: Node 0 CPU frequency 1200 MHz
RTT Norm:44
NBU0 DRAM BAR0 base: 00000000 limit: 0013f000 xlate: 00000001 node: 00000000
( 0 MB -> 320 MB
, size: 320 MB)
NBU0 DRAM BAR1 base: 001d0000 limit: 0088f000 xlate: 00090001 node: 00000000
```

```
( 464 MB -> 2192 MB
, size: 1728 MB)
Modifying Default Flash Address map..Done
Initialized eMMC Host Controller
Detected SD Card
BLC is 1 (preset 10)
Hit any key to stop autoboot: 0
Boot Image selection
Reading the Boot Block Info...Passed !!
Images are OK A:0x0 B:0x0
Boot Selector set to Bootflash Partition A image...
Verifying Copyright Information..success for Image - 0
Boot Selector: Booting Bootflash Partition A image...
Copying stage-2 loader from 0xb6120000 to 0x8c100000(size = 0x100000)
Boot Image selection DONE.
## Starting application at 0x8C100000 ...
```

```
U-Boot 2010.03-rcl(Dell Force10)
Built by build at tools-sjc-01 on Thu May 31 23:53:38 2012
IOM Boot Label 4.0.1.0
```

```
DRAM: 2 GB
Initialized CPLD on CS3
Detected [XLP308 (Lite+) Rev A0]
Initializing I2C0: speed = 30 KHz, prescaler = 0x0377 -- done.
Initializing I2C1: speed = 100 KHz, prescaler = 0x0109 -- done.
Initialized eMMC Host Controller
Detected SD Card
Now running in RAM - U-Boot [N64 ABI, Big-Endian] at: ffffffff8c100000
Flash: 256 MB
PCIE (B0:D01:F0) : Link up.
PCIE (B0:D01:F1) : No Link.
In: serial
Out: serial
Err: serial
Net: nae-0: PHY is Broadcom BCM54616S
```

--More--

```
SOFTWARE IMAGE HEADER DATA :
-----
```

--More--

Starting Dell Networking application

Welcome to Dell Easy Setup Wizard

The setup wizard guides you through the initial switch configuration, and gets you up and running as quickly as possible. You can skip the setup wizard, and enter CLI mode to manually configure the switch. You must respond to the next question to run the setup wizard within 60 seconds, otherwise the system will continue with normal operation using the default system configuration.
Note: You can exit the setup wizard at any point by entering [ctrl+c].

```
Would you like to run the setup wizard (you must answer this question within
60 seconds)? [Y/N]: N
00:00:40: %STKUNIT0-M:CP %IFMGR-5-ASTATE_UP: Changed interface Admin state to
up: Vl 1
00:00:42: %STKUNIT0-M:CP %IFMGR-5-IFM_ISCSI_ENABLE: iSCSI has been enabled
causing flow control
to be enabled on all interfaces.
```

```
EQL detection and enabling iscsi profile-compellent on an interface may cause
some automatic
configurations to occur like jumbo frames on all ports and no storm control
and spanning tree port-fast on the port of detection
00:00:42: %STKUNIT0-M:CP %SEC-5-LOGIN_SUCCESS: Login successful for user on
line console
Dell>en
Password:
```

Default Configuration

A version of the Dell Networking OS is pre-loaded onto the chassis; however, the system is not configured when you power up for the first time (except for the default hostname, which is `Dell`). You must configure the system using the CLI.

Configuring a Host Name

The host name appears in the prompt. The default host name is `Dell`.

- Host names must start with a letter and end with a letter or digit.
- Characters within the string can be letters, digits, and hyphens.

To create a host name, use the following command.

- Create a host name.
CONFIGURATION mode

`hostname name`

Example of the `hostname` Command

```
Dell(conf)#hostname R1
R1(conf)#
```

Configuring a Host Name

The host name appears in the prompt. The default host name is `Dell`.

- Host names must start with a letter and end with a letter or digit.
- Characters within the string can be letters, digits, and hyphens.

To create a host name, use the following command.

- Create a host name.
CONFIGURATION mode

`hostname name`

Example of the `hostname` Command

```
Dell(conf)#hostname R1
R1(conf)#
```


Accessing the System Remotely

You can configure the system to access it remotely by Telnet or SSH.

The MXL 10/40GbE switch IO module has a dedicated management port and a management routing table that is separate from the IP routing table.

Accessing the MXL Switch Remotely

Configuring the system for Telnet is a three-step process, as described in the following topics:

1. Configure an IP address for the management port. [Configure the Management Port IP Address](#)
2. Configure a management route with a default gateway. [Configure a Management Route](#)
3. Configure a username and password. [Configure a Username and Password](#)

Configure the Management Port IP Address

To access the system remotely, assign IP addresses to the management ports.

1. Enter INTERFACE mode for the Management port.
CONFIGURATION mode

```
interface ManagementEthernet slot/port
```

- *slot*: the range is 0.
- *port*: the range is 0.

2. Assign an IP address to the interface.
INTERFACE mode

```
ip address ip-address/mask
```

- *ip-address*: an address in dotted-decimal format (A.B.C.D).
- *mask*: a subnet mask in /prefix-length format (/ xx).

3. Enable the interface.
INTERFACE mode

```
no shutdown
```

Configure a Management Route

Define a path from the system to the network from which you are accessing the system remotely. Management routes are separate from IP routes and are only used to manage the system through the management port.

To configure a management route, use the following command.

- Configure a management route to the network from which you are accessing the system.
CONFIGURATION mode

```
management route ip-address/mask gateway
```

- *ip-address*: the network address in dotted-decimal format (A.B.C.D).
- *mask*: a subnet mask in /prefix-length format (/ xx).
- *gateway*: the next hop for network traffic originating from the management port.

Configuring a Username and Password

To access the system remotely, configure a system username and password. To configure a system username and password, use the following command.

- Configure a username and password to access the system remotely.
CONFIGURATION mode

```
username username password [encryption-type] password
```

- *encryption-type*: specifies how you are inputting the password, is 0 by default, and is not required.
 - * 0 is for inputting the password in clear text.
 - * 7 is for inputting a password that is already encrypted using a Type 7 hash. Obtaining the encrypted password from the configuration of another Dell Networking system.

Configuring the Enable Password

Access EXEC Privilege mode using the `enable` command. EXEC Privilege mode is unrestricted by default. Configure a password as a basic security measure.

There are two types of enable passwords:

- `enable password` stores the password in the running/startup configuration using a DES encryption method.
- `enable secret` is stored in the running/startup configuration in using a stronger, MD5 encryption method.

Dell Networking recommends using the `enable secret` password.

To configure an enable password, use the following command.

- Create a password to access EXEC Privilege mode.
CONFIGURATION mode


```
enable [password | secret] [level level] [encryption-type] password
```

- *level*: is the privilege level, is 15 by default, and is not required
- *encryption-type*: specifies how you are inputting the password, is 0 by default, and is not required.
 - * 0 is for inputting the password in clear text.
 - * 7 is for inputting a password that is already encrypted using a DES hash. Obtain the encrypted password from the configuration file of another Dell Networking system. You can only use this for the `enable password`.

- * 5 is for inputting a password that is already encrypted using an MD5 hash. Obtain the encrypted password from the configuration file of another Dell Networking system. You can only use this for the `enable secret` password.


Configuration File Management

Files can be stored on and accessed from various storage media. Rename, delete, and copy files on the system from EXEC Privilege mode.

 **NOTE:** Using flash memory cards in the system that have not been approved by Dell Networking can cause unexpected system behavior, including a reboot.

Copy Files to and from the System

The command syntax for copying files is similar to UNIX. The copy command uses the format `copy source-file-url destination-file-url`.

 **NOTE:** For a detailed description of the copy command, refer to the *Dell Networking OS Command Line Reference Guide*.


- To copy a local file to a remote system, combine the `file-origin` syntax for a local file location with the `file-destination` syntax for a remote file location.
- To copy a remote file to Dell Networking system, combine the `file-origin` syntax for a remote file location with the `file-destination` syntax for a local file location.

Table 3. Forming a copy Command

Location	<i>source-file-url</i> Syntax	<i>destination-file-url</i> Syntax
Internal flash: flash	<code>copy flash://filename</code>	<code>flash://filename</code>
USB flash: usbflash	<code>usbflash://filename</code>	<code>usbflash://filename</code>
For a remote file location: FTP server	<code>copy ftp:// username:password@{hostip hostname}/filepath/ filename</code>	<code>ftp:// username:password@{hostip hostname}/ filepath/ filename</code>
For a remote file location: TFTP server	<code>copy tftp://{hostip hostname}/filepath/ filename</code>	<code>tftp://{hostip hostname}/filepath/ filename</code>
For a remote file location: SCP server	<code>copy scp://{hostip hostname}/filepath/ filename</code>	<code>scp://{hostip hostname}/filepath/ filename</code>

Important Points to Remember

- You may not copy a file from one remote system to another.
- You may not copy a file from one location to the same location.
- When copying to a server, you can only use a hostname if you configured a domain name server (DNS) server.

 **NOTE:** If all of the following conditions are true, the Portmode Hybrid configuration is not applied, because of the configuration process for server ports as switch ports by default:

- The running configuration is saved in flash.
- The startup configuration is deleted.
- The switch is reloaded.
- The saved configuration is copied to the running configuration.

To avoid this scenario, delete the switch port configuration from the running configuration before copying the saved configuration to the running configuration.

Example of Copying a File to an FTP Server

Example of Importing a File to the Local System

The bold flash shows the local location and the bold ftp shows the remote location.

```
Dell#copy flash:://FTOS-EF-8.2.1.0.bin ftp:://myusername:mypassword@10.10.10.10/  
/FTOS/FTOS-EF-8.2.1.0  
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!  
27952672 bytes successfully copied
```

```
core1#$//copy ftp:://myusername:mypassword@10.10.10.10//FTOS/  
FTOS-EF-8.2.1.0.bin flash::/  
Destination file name [FTOS-EF-8.2.1.0.bin.bin]:  
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!  
26292881 bytes successfully copied
```

Save the Running-Configuration

The running-configuration contains the current system configuration. Dell Networking recommends copying your running-configuration to the startup-configuration.

The system uses the startup-configuration during boot-up to configure the system. The startup-configuration is stored in the internal flash on the IOM by default, but you can save it to a USB flash device or a remote server.

The commands in this section follow the same format as those commands in the [Copy Files to and from the System](#) section but use the filenames *startup-config* and *running-config*. These commands assume that current directory is the internal flash, which is the system default.

- Save the running-config to the startup-configuration on the internal flash.
EXEC Privilege mode

```
copy running-config startup-config
```

- Save the running-configuration on the IOM.
EXEC Privilege mode

```
copy running-config usbflash://filename
```

- Save the running-configuration to an FTP server.
EXEC Privilege mode

```
copy running-config ftp:// username:password@{hostip | hostname}/filepath/  
filename
```

- Save the running-configuration to a TFTP server.

EXEC Privilege mode

```
copy running-config tftp://{hostip | hostname}/ filepath/filename
```

- Save the running-configuration to an SCP server.

EXEC Privilege mode

```
copy running-config scp://{hostip | hostname}/ filepath/filename
```



NOTE: When copying to a server, you can only use a host name if you have configured a DNS server.

- Save the running-configuration to the startup-configuration on the internal flash of the primary RPM. Then copy the new startup-config file to the external flash of the primary RPM.

EXEC Privilege mode

```
copy running-config startup-config duplicate
```

Dell Networking OS Behavior: If you create a startup-configuration on an RPM and then move the RPM to another chassis, the startup-configuration is stored as a backup file (with the extension *.bak*), and a new, empty startup-configuration file is created. To restore your original startup-configuration in this situation, overwrite the new startup-configuration with the original one using the `copy startup-config.bak startup-config` command.

Viewing Files

You can only view file information and content on local file systems.

To view a list of files or the contents of a file, use the following commands.

- View a list of files on the internal flash.

EXEC Privilege mode

```
dir flash:
```

- View a list of files on the usbflash.

EXEC Privilege mode

```
dir usbflash:
```

- View the contents of a file in the internal flash.

EXEC Privilege mode

```
show file flash://filename
```

- View the contents of a file in the usb flash.

EXEC Privilege mode

```
show file usbflash://filename
```

- View the running-configuration.

EXEC Privilege mode

```
show running-config
```

- View the startup-configuration.

EXEC Privilege mode

```
show startup-config
```

Example of the dir Command

The output of the `dir` command also shows the read/write privileges, size (in bytes), and date of modification for each file.

```
Dell#dir
Directory of flash:

 1 drwx      4096 Jan  1 1980 00:00:00 +00:00 .
 2 drwx      2048 May 10 2011 14:45:15 +00:00 ..
 3 drwx      4096 Feb 17 2011 00:28:00 +00:00 TRACE_LOG_DIR
 4 drwx      4096 Feb 17 2011 00:28:02 +00:00 CORE_DUMP_DIR
 5 d---      4096 Feb 17 2011 00:28:02 +00:00 ADMIN_DIR
 6 -rwx       1272 Apr 29 2011 16:15:14 +00:00 startup-config
 7 -rwx     10093 Feb 17 2011 20:48:02 +00:00 abhi-jan26.cfg
 8 -rwx    217155 Feb 22 2011 23:14:34 +00:00 show-tech-cfg.txt
 9 -rwx       5162 Mar  2 2011 04:02:58 +00:00 runn-feb6
10 -rwx     10507 Mar  3 2011 01:17:16 +00:00 abhi-feb7.cfg
11 -rwx         4 May  6 2011 22:05:06 +00:00 dhcpBindConflict
12 -rwx       6900 Feb 17 2011 04:43:12 +00:00 startup-config.bak
13 -rwx   1244038 Feb 13 2011 04:27:16 +00:00 f10cp_sysd_110213042625.acore.gz

flash: 2143281152 bytes total (2123755520 bytes free)
--More--
```

View Configuration Files

Configuration files have three commented lines at the beginning of the file, as shown in the following example, to help you track the last time any user made a change to the file, which user made the changes, and when the file was last saved to the startup-configuration.

In the running-configuration file, if there is a difference between the timestamp on the “Last configuration change,” and “Startup-config last updated,” you have made changes that have not been saved and will not be preserved after a system reboot.

Example of the show running-config Command

```
Dell#show running-config
Current Configuration ...
Current Configuration ...
! Version E8-3-16-0
! Last configuration change at Tue Mar 6 11:51:50 2012 by default
! Startup-config last updated at Tue Mar 6 07:41:23 2012 by default
!
boot system stack-unit 5 primary tftp://10.11.200.241/dt-m1000e-3-a2
boot system stack-unit 5 secondary system: B:
boot system stack-unit 5 default tftp://10.11.200.241/dt-m1000e-3-b2
boot system gateway 10.11.209.254
--More--
```

Managing the File System

The Dell Networking system can use the internal Flash, USB Flash, or remote devices to store files. The system stores files on the internal Flash by default but you can configure the system to store files elsewhere.

To view file system information, use the following command.

- View information about each file system.
EXEC Privilege mode

```
show file-systems
```

The output of the `show file-systems` command in the following example shows the total capacity, amount of free memory, file structure, media type, read/write privileges for each storage device in use.

```
Dell#show file-systems
Size(b)      Free(b)      Feature Type      Flags Prefixes
2143281152  2000785408  FAT32  USERFLASH  rw  flash:
15848660992  831594496   FAT32  USBFLASH   rw  usbflash:
-           -           -      network   rw  ftp:
-           -           -      network   rw  tftp:
-           -           -      network   rw  scp:
```

You can change the default file system so that file management commands apply to a particular device or memory.

To change the default directory, use the following command.

- Change the default directory.
EXEC Privilege mode

```
cd directory
```

You can change the default storage location to the USB Flash, as shown. File management commands then apply to the USB Flash rather than the internal Flash. The bold lines show that no file system is specified and that the file is saved to an USB Flash.

```
Dell#cd usbflash:
Dell#copy running-config test
!
3998 bytes successfully copied

DellS#dir
Directory of usbflash:

 1 drwx 4096 Jan 01 1980 00:00:00 +00:00 .
 2 drwx 2048 May 02 2012 07:05:06 +00:00 ..
 3 -rwx 1272 Apr 29 2011 16:15:14 +00:00 startup-config
 4 -rwx 3998 May 11 2011 23:36:12 +00:00 test
```

View the Command History

The command-history trace feature captures all commands entered by all users of the system with a time stamp and writes these messages to a dedicated trace log buffer.

The system generates a trace message for each executed command. No password information is saved to the file.

To view the command-history trace, use the `show command-history` command.

Example of the `show command-history` Command


```
Dell#show command-history
[5/18 21:58:32]: CMD-(TEL0):[enable]by admin from vty0 (10.11.68.5)
[5/18 21:58:48]: CMD-(TEL0):[configure]by admin from vty0 (10.11.68.5)
    - Repeated 1 time.
[5/18 21:58:57]: CMD-(TEL0):[interface port-channel 1]by admin from vty0
(10.11.68.5)
[5/18 21:59:9]: CMD-(TEL0):[show config]by admin from vty0 (10.11.68.5)
[5/18 22:4:32]: CMD-(TEL0):[exit]by admin from vty0 (10.11.68.5)
[5/18 22:4:41]: CMD-(TEL0):[show interfaces port-channel brief]by admin from
vty0
(10.11.68.5)
```

Using HTTP for File Transfers

Starting with Release 9.3(0.1), you can use HTTP to copy files or configuration details to a remote server. Use the **copy source-file-url http://host[:port]/file-path** command to transfer files to an external server. Enter the following *source-file-url* keywords and information:

- To copy a file from the internal FLASH, enter `flash://` followed by the filename.
- To copy the running configuration, enter the keyword `running-config`.
- To copy the startup configuration, enter the keyword `startup-config`.
- To copy a file on the external FLASH, enter `usbflash://` followed by the filename.

Upgrading and Downgrading the Dell Networking OS

 **NOTE:** To upgrade the Dell Networking OS, refer to the Release Notes for the version you want to load on the system.

Using Hashes to Validate Software Images

You can use the MD5 message-digest algorithm or SHA256 Secure Hash Algorithm to validate the software image on the flash drive, after the image has been transferred to the system, but before the image has been installed. The validation calculates a hash value of the downloaded image file on system's flash drive, and, optionally, compares it to a Dell Networking published hash for that file.

The MD5 or SHA256 hash provides a method of validating that you have downloaded the original software. Calculating the hash on the local image file, and comparing the result to the hash published for that file on iSupport, provides a high level of confidence that the local copy is exactly the same as the published software image. This validation procedure, and the **verify {md5 | sha256}** command to support it, can prevent the installation of corrupted or modified images.

The **verify {md5 | sha256}** command calculates and displays the hash of any file on the specified local flash drive. You can compare the displayed hash against the appropriate hash published on i-Support. Optionally, the published hash can be included in the **verify {md5 | sha256}** command, which will display whether it matches the calculated hash of the indicated file.

To validate a software image:

1. Download Dell Networking OS software image file from the iSupport page to the local (FTP or TFTP) server. The published hash for that file is displayed next to the software image file on the iSupport page.
2. Go on to the Dell Networking system and copy the software image to the flash drive, using the **copy** command.
3. Run the **verify {md5 | sha256} [flash://img-file [hash-value]** command. For example, **verify sha256 flash://FTOS-SE-9.5.0.0.bin**
4. Compare the generated hash value to the expected hash value published on the iSupport page.

To validate the software image on the flash drive after the image has been transferred to the system, but before the image has been installed, use the **verify {md5 | sha256} [flash://img-file [hash-value]** command in EXEC mode.

- **md5**: MD5 message-digest algorithm
- **sha256**: SHA256 Secure Hash Algorithm
- **flash**: (Optional) Specifies the flash drive. The default is to use the flash drive. You can just enter the image file name.
- **hash-value**: (Optional). Specify the relevant hash published on i-Support.
- **img-file**: Enter the name **of** the Dell Networking **software** image file to validate

Examples: Without Entering the Hash Value for Verification

MD5

```
Dell# verify md5 flash://FTOS-SE-9.5.0.0.bin
MD5 hash for FTOS-SE-9.5.0.0.bin: 275ceb73a4f3118e1d6bcf7d75753459
```

SHA256

```
Dell# verify sha256 flash://FTOS-SE-9.5.0.0.bin
SHA256 hash for FTOS-SE-9.5.0.0.bin:
e6328c06faf814e6899ceead219afbf9360e986d692988023b749e6b2093e933
```

Examples: Entering the Hash Value for Verification

MD5

```
Dell# verify md5 flash://FTOS-SE-9.5.0.0.bin 275ceb73a4f3118e1d6bcf7d75753459
MD5 hash VERIFIED for FTOS-SE-9.5.0.0.bin
```

SHA256

```
Dell# verify sha256 flash://FTOS-SE-9.5.0.0.bin
e6328c06faf814e6899ceead219afbf9360e986d692988023b749e6b2093e933
SHA256 hash VERIFIED for FTOS-SE-9.5.0.0.bin
```

Management

Management is supported on the Dell Networking MXL 10/40GbE Switch IO Module.

This chapter describes the different protocols or services used to manage the Dell Networking system.

Configuring Privilege Levels

Privilege levels restrict access to commands based on user or terminal line.

There are 15 privilege levels, of which two are pre-defined. The default privilege level is **1**.

- **Level 1** — Access to the system begins at EXEC mode, and EXEC mode commands are limited to basic commands, some of which are `enable`, `disable`, and `exit`.
- **Level 15** — To access all commands, enter EXEC Privilege mode. Normally, enter a password to enter this mode.

Creating a Custom Privilege Level

Custom privilege levels start with the default EXEC mode command set.

You can then customize privilege levels 2-14 by:

- removing commands from the EXEC mode commands
- moving commands from EXEC Privilege mode to EXEC mode
- allowing access to CONFIGURATION mode commands
- allowing access to INTERFACE, LINE, ROUTE-MAP, and ROUTER mode commands

You can access all commands at your privilege level and below.

Moving a Command from EXEC Privilege Mode to EXEC Mode

Remove a command from the list of available commands in EXEC mode for a specific privilege level using the `privilege exec` command from CONFIGURATION mode. In the command, specify a level *greater* than the level given to a user or terminal line, then the first keyword of each restricted command.

Moving a Command from EXEC Privilege Mode to EXEC Mode

Move a command from EXEC Privilege to EXEC mode for a privilege level using the `privilege exec` command from CONFIGURATION mode. In the command, specify the privilege level of the user or terminal line, and specify *all* keywords in the command to which you want to allow access.

Allowing Access to CONFIGURATION Mode Commands

Allow access to CONFIGURATION mode using the `privilege exec level level` command configure from CONFIGURATION mode. A user that enters CONFIGURATION mode remains at his

privilege level, and has access to only two commands, end and exit. Individually specify each CONFIGURATION mode command to which you want to allow access using the `privilege configure level level` command. In the command, specify the privilege level of the user or terminal line, and specify *all* keywords in the command to which you want to allow access.

Allowing Access to INTERFACE, LINE, ROUTE-MAP, and ROUTER Mode

1. Similar to allowing access to CONFIGURATION mode, to allow access to INTERFACE, LINE, ROUTE-MAP, and ROUTER modes, first allow access to the command that enters you into the mode. For example, allow a user to enter INTERFACE mode using the `privilege configure level level interface gigabitethernet` command.
2. Then, individually identify the INTERFACE, LINE, ROUTE-MAP or ROUTER commands to which you want to allow access using the `privilege {interface | line | route-map | router} level level` command. In the command, specify the privilege level of the user or terminal line and specify all keywords in the command to which you want to allow access.

Customizing a Privilege Level

to customize a privilege level, use the following commands.

1. Remove a command from the list of available commands in EXEC mode.
CONFIGURATION mode

```
privilege exec level level {command ||...|| command}
```

2. Move a command from EXEC Privilege to EXEC mode.
CONFIGURATION mode

```
privilege exec level level {command ||...|| command}
```

3. Allow access to CONFIGURATION mode.
CONFIGURATION mode

```
privilege exec configure level level
```

4. Allow access to INTERFACE, LINE, ROUTE-MAP, and/or ROUTER mode. Specify all keywords in the command.
CONFIGURATION mode

```
privilege configure level level {interface | line | route-map | router}  
{command-keyword ||...|| command-keyword}
```

5. Allow access to a CONFIGURATION, INTERFACE, LINE, ROUTE-MAP, and/or ROUTER mode command.
CONFIGURATION mode

```
privilege {configure | interface | line | route-map | router} level level  
{command ||...|| command}
```

Create a Custom Privilege Level Apply a Privilege Level to a Username

The following configuration privilege level 3. This level:

- removes the `resequence` command from EXEC mode by requiring a minimum of privilege level 4
- moves the `capture bgp-pdu max-buffer-size` command from EXEC Privilege to EXEC mode by requiring a minimum privilege level 3, which is the configured level for VTY 0

- allows access to CONFIGURATION mode with the banner command
- allows access to INTERFACE and LINE modes with the no command

```

Dell(conf)#do show run privilege
!
Dell(conf)#privilege exec level 3 capture
Dell(conf)#privilege exec level 3 configure
Dell(conf)#privilege exec level 4 resequence
Dell(conf)#privilege exec level 3 clear arp-cache
Dell(conf)#privilege exec level 3 clear arp-cache max-buffer-size
Dell(conf)#privilege configure level 3 line
Dell(conf)#privilege configure level 3 interface
Dell(conf)#do telnet 10.11.80.201
[telnet output omitted]
Dell#show priv
Current privilege level is 3.
Dell#?
capture          Capture packet
configure        Configuring from terminal
disable         Turn off privileged commands
enable          Turn on privileged commands
exit            Exit from the EXEC
ip              Global IP subcommands
monitor         Monitoring feature
mtrace          Trace reverse multicast path from destination to source
ping            Send echo messages
quit            Exit from the EXEC
show            Show running system information
[output omitted]
Dell#config
[output omitted]
Dell(conf)#do show priv
Current privilege level is 3.
Dell(conf)#?
end              Exit from configuration mode
exit            Exit from configuration mode
interface       Select an interface to configure
Dell(conf)#interface ?
loopback        Loopback interface
managementethernet Management Ethernet interface
null            Null interface
port-channel    Port-channel interface
range           Configure interface range
tengigabitethernet TenGigabit Ethernet interface
vlan            VLAN interface
Dell(conf)#interface tengigabitethernet 1/1
Dell(conf-if-te-1/1)#?
end              Exit from configuration mode
exit            Exit from interface configuration mode
Dell(conf-if-te-1/1)#exit
Dell(conf)#line ?
console         Primary terminal line
vty             Virtual terminal
Dell(conf)#line vty 0
Dell(conf-line-vty)#?
exit            Exit from line configuration mode
Dell(conf-line-vty)#

```

Applying a Privilege Level to a Username

To set the user privilege level, use the following command.

- Configure a privilege level for a user.

CONFIGURATION mode

```
username username privilege level
```


Applying a Privilege Level to a Terminal Line

To set a privilege level for a terminal line, use the following command.

- Configure a privilege level for a terminal line.

Line mode

```
privilege levellevel
```

 **NOTE:** When you assign a privilege level between 2 and 15, access to the system begins at EXEC mode, but the prompt is `hostname#`, rather than `hostname>`.

Configuring Logging

The Dell Networking operating system tracks changes in the system using event and error messages. By default, the system logs these messages on:

- the internal buffer
- console and terminal lines
- any configured syslog servers

To disable logging, use the following commands.

- Disable all logging except on the console.

CONFIGURATION mode

```
no logging on
```

- Disable logging to the logging buffer.

CONFIGURATION mode

```
no logging buffer
```

- Disable logging to terminal lines.

CONFIGURATION mode

```
no logging monitor
```

- Disable console logging.

CONFIGURATION mode

```
no logging console
```

Audit and Security Logs

This section describes how to configure, display, and clear audit and security logs.

The following is the configuration task list for audit and security logs:

- [Enabling Audit and Security Logs](#)
- [Displaying Audit and Security Logs](#)
- [Clearing Audit Logs](#)

Enabling Audit and Security Logs

You enable audit and security logs to monitor configuration changes or determine if these changes affect the operation of the system in the network. You log audit and security events to a system log server, using the **logging extended** command in CONFIGURATION mode.

Audit Logs

The audit log contains configuration events and information. The types of information in this log consist of the following:

- User logins to the switch.
- System events for network issues or system issues.
- Users making configuration changes. The switch logs who made the configuration changes and the date and time of the change. However, each specific change on the configuration is not logged. Only that the configuration was modified is logged with the user ID, date, and time of the change.
- Uncontrolled shutdown.

Security Logs


The security log contains security events and information. RBAC restricts access to audit and security logs based on the CLI sessions' user roles. The types of information in this log consist of the following:

- Establishment of secure traffic flows, such as SSH.
- Violations on secure flows or certificate issues.
- Adding and deleting of users.
- User access and configuration changes to the security and crypto parameters (not the key information but the crypto configuration)

Important Points to Remember

When you enabled RBAC and extended logging:

- Only the system administrator user role can execute this command.
- The system administrator and system security administrator user roles can view security events and system events.
- The system administrator user roles can view audit, security, and system events.
- Only the system administrator and security administrator user roles can view security logs.
- The network administrator and network operator user roles can view system events.

 **NOTE:** If extended logging is disabled, you can only view system events, regardless of RBAC user role.

Example of Enabling Audit and Security Logs

```
Dell(conf)#logging extended
```

Displaying Audit and Security Logs

To display audit logs, use the `show logging auditlog` command in Exec mode. To view these logs, you must first enable the logging extended command. Only the RBAC system administrator user role can view the audit logs. Only the RBAC security administrator and system administrator user role can view the security logs. If extended logging is disabled, you can only view system events, regardless of RBAC user role. To view security logs, use the `show logging` command.

Example of the `show logging auditlog` Command

For information about the logging extended command, see [Enabling Audit and Security Logs](#)

```
Dell#show logging auditlog
May 12 12:20:25: Dell#: %CLI-6-logging extended by admin from vty0 (10.14.1.98)
May 12 12:20:42: Dell#: %CLI-6-configure terminal by admin from vty0
(10.14.1.98)
May 12 12:20:42: Dell#: %CLI-6-service timestamps log datetime by admin from
vty0 (10.14.1.98)
```

Example of the `show logging` Command for Security

For information about the logging extended command, see [Enabling Audit and Security Logs](#)

```
Dell#show logging
Jun 10 04:23:40: %STKUNIT0-M:CP %SEC-5-LOGIN_SUCCESS: Login successful for user
admin on line vty0 ( 10.14.1.91 )
```

Clearing Audit Logs

To clear audit logs, use the `clear logging auditlog` command in Exec mode. When RBAC is enabled, only the system administrator user role can issue this command.

Example of the `clear logging auditlog` Command

```
Dell# clear logging auditlog
```

Configuring Logging Format

To display syslog messages in a RFC 3164 or RFC 5424 format, use the `logging version [0 | 1]` command in CONFIGURATION mode. By default, the system log version is set to 0.

The following describes the two log messages formats:

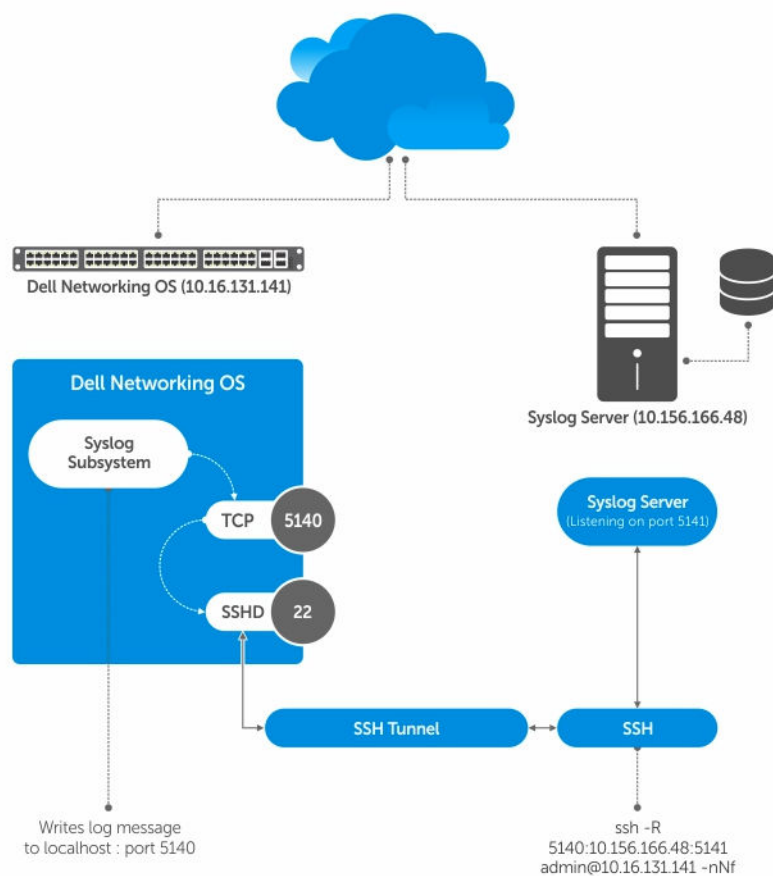
- **0** – Displays syslog messages format as described in RFC 3164, The BSD syslog Protocol
- **1** – Displays syslog message format as described in RFC 5424, The SYSLOG Protocol

Example of Configuring the Logging Message Format

```
Dell(conf)#logging version ?
<0-1> Select syslog version (default = 0)
Dell(conf)#logging version 1
```

Setting Up a Secure Connection to a Syslog Server

You can use reverse tunneling with the port forwarding to securely connect to a syslog server.



Pre-requisites

To configure a secure connection from the switch to the syslog server:

1. On the switch, enable the SSH server
`Dell(conf)#ip ssh server enable`
2. On the syslog server, create a reverse SSH tunnel from the syslog server to FTOS switch, using following syntax:

```
ssh -R <remote port>:<syslog server>:<syslog server listen port>  
user@remote_host -nNf
```


In the following example the syslog server IP address is 10.156.166.48 and the listening port is 5141. The switch IP address is 10.16.131.141 and the listening port is 5140

```
ssh -R 5140:10.156.166.48:5141 admin@10.16.131.141 -nNf
```

3. Configure logging to a local host. *localhost* is "127.0.0.1" or "::1".

If you do not, the system displays an error when you attempt to enable role-based only AAA authorization.

```
Dell(conf)# logging localhost tcp port
Dell(conf)#logging 127.0.0.1 tcp 5140
```

Display the Logging Buffer and the Logging Configuration

To display the current contents of the logging buffer and the logging settings for the system, use the `show logging` command in EXEC privilege mode. When RBAC is enabled, the security logs are filtered based on the user roles. Only the security administrator and system administrator can view the security logs.

Example of the `show logging` Command

```
Dell#show logging
syslog logging: enabled
  Console logging: level Debugging
  Monitor logging: level Debugging
  Buffer logging: level Debugging, 40 Messages Logged, Size (40960 bytes)
  Trap logging: level Informational
%IRC-6-IRC_COMMUP: Link to peer RPM is up
%RAM-6-RAM_TASK: RPM1 is transitioning to Primary RPM.
%RPM-2-MSG:CP1 %POLLGR-2-MMC_STATE: External flash disk missing in 'slot0:'
%CHMGR-5-CARDDETECTED: Line card 0 present
%CHMGR-5-CARDDETECTED: Line card 2 present
%CHMGR-5-CARDDETECTED: Line card 4 present
%CHMGR-5-CARDDETECTED: Line card 5 present
%CHMGR-5-CARDDETECTED: Line card 8 present
%CHMGR-5-CARDDETECTED: Line card 10 present
%CHMGR-5-CARDDETECTED: Line card 12 present
%TSM-6-SFM_DISCOVERY: Found SFM 0
%TSM-6-SFM_DISCOVERY: Found SFM 1
%TSM-6-SFM_DISCOVERY: Found SFM 2
%TSM-6-SFM_DISCOVERY: Found SFM 3
%TSM-6-SFM_DISCOVERY: Found SFM 4
%TSM-6-SFM_DISCOVERY: Found SFM 5
%TSM-6-SFM_DISCOVERY: Found SFM 6
%TSM-6-SFM_DISCOVERY: Found SFM 7
%TSM-6-SFM_SWITCHFAB_STATE: Switch Fabric: UP
%TSM-6-SFM_DISCOVERY: Found SFM 8
%TSM-6-SFM_DISCOVERY: Found 9 SFMs
%CHMGR-5-CHECKIN: Checkin from line card 5 (type EX1YB, 1 ports)
%TSM-6-PORT_CONFIG: Port link status for LC 5 => portpipe 0: OK portpipe 1: N/A
%CHMGR-5-LINECARDUP: Line card 5 is up
%CHMGR-5-CHECKIN: Checkin from line card 12 (type S12YC12, 12 ports)
%TSM-6-PORT_CONFIG: Port link status for LC 12 => portpipe 0: OK portpipe 1: N/A
%CHMGR-5-LINECARDUP: Line card 12 is up
%IFMGR-5-CSTATE_UP: changed interface Physical state to up: So 12/8
%IFMGR-5-CSTATE_DN: changed interface Physical state to down: So 12/8
```

To view any changes made, use the `show running-config logging` command in EXEC privilege mode.

Log Messages in the Internal Buffer

All error messages, except those beginning with `%BOOTUP (Message)`, are log in the internal buffer. For example, `%BOOTUP:RPM0:CP %PORTPIPE-INIT-SUCCESS: Portpipe 0 enabled`

Configuration Task List for System Log Management

There are two configuration tasks for system log management:

- [Disabling System Logging](#)
- [Sending System Messages to a Syslog Server](#)

Disabling System Logging

By default, logging is enabled and log messages are sent to the logging buffer, all terminal lines, the console, and the syslog servers.

To disable system logging, use the following commands.

- Disable all logging except on the console.

CONFIGURATION mode

```
no logging on
```

- Disable logging to the logging buffer.

CONFIGURATION mode

```
no logging buffer
```

- Disable logging to terminal lines.

CONFIGURATION mode

```
no logging monitor
```

- Disable console logging.

CONFIGURATION mode

```
no logging console
```

Sending System Messages to a Syslog Server

To send system messages to a specified syslog server, use the following command. The following syslog standards are supported: RFC 5424 The SYSLOG Protocol, R. Gerhards and Adiscon GmbH, March 2009, obsoletes RFC 3164 and RFC 5426 Transmission of Syslog messages over UDP.

- Specify the server to which you want to send system messages. You can configure up to eight syslog servers.

CONFIGURATION mode

```
logging {ip-address | ipv6-address | hostname} {{udp {port}} | {tcp {port}}}
```

Configuring a UNIX System as a Syslog Server

To configure a UNIX System as a syslog server, use the following command.

- Configure a UNIX system as a syslog server by adding the following lines to `/etc/syslog.conf` on the UNIX system and assigning write permissions to the file.
 - Add line on a 4.1 BSD UNIX system. `local7.debugging /var/log/log7.log`
 - Add line on a 5.7 SunOS UNIX system. `local7.debugging /var/adm/ftos.log`

In the previous lines, `local7` is the logging facility level and `debugging` is the severity level.

Changing System Logging Settings

You can change the default settings of the system logging by changing the severity level and the storage location.

The default is to log all messages up to debug level, that is, all system messages. By changing the severity level in the logging commands, you control the number of system messages logged.

To specify the system logging settings, use the following commands.

- Specify the minimum severity level for logging to the logging buffer.

CONFIGURATION mode

```
logging buffered level
```

- Specify the minimum severity level for logging to the console.

CONFIGURATION mode

```
logging console level
```

- Specify the minimum severity level for logging to terminal lines.

CONFIGURATION mode

```
logging monitor level
```

- Specify the minimum severity level for logging to a syslog server.

CONFIGURATION mode

```
logging trap level
```

- Specify the minimum severity level for logging to the syslog history table.


CONFIGURATION mode

```
logging history level
```

- Specify the size of the logging buffer.

CONFIGURATION mode

```
logging buffered size
```

 **NOTE:** When you decrease the buffer size, the system deletes all messages stored in the buffer. Increasing the buffer size does not affect messages in the buffer.

- Specify the number of messages that the system saves to its logging history table.

CONFIGURATION mode

```
logging history size size
```

To view the logging buffer and configuration, use the `show logging` command in EXEC privilege mode, as shown in the example for [Display the Logging Buffer and the Logging Configuration](#).

To view the logging configuration, use the `show running-config logging` command in privilege mode, as shown in the example for [Configuring a UNIX Logging Facility Level](#).

Display the Logging Buffer and the Logging Configuration

To display the current contents of the logging buffer and the logging settings for the system, use the `show logging` command in EXEC privilege mode. When RBAC is enabled, the security logs are filtered based on the user roles. Only the security administrator and the system administrator can view the security logs.

Example of the `show logging` Command

```
Dell#show logging
syslog logging: enabled
  Console logging: level Debugging
  Monitor logging: level Debugging
  Buffer logging: level Debugging, 40 Messages Logged, Size (40960 bytes)
  Trap logging: level Informational
%IRC-6-IRC_COMMUP: Link to peer RPM is up
%RAM-6-RAM_TASK: RPM1 is transitioning to Primary RPM.
%RPM-2-MSG:CP1 %POLLGR-2-MMC_STATE: External flash disk missing in 'slot0:'
%CHMGR-5-CARDDETECTED: Line card 0 present
%CHMGR-5-CARDDETECTED: Line card 2 present
%CHMGR-5-CARDDETECTED: Line card 4 present
%CHMGR-5-CARDDETECTED: Line card 5 present
%CHMGR-5-CARDDETECTED: Line card 8 present
%CHMGR-5-CARDDETECTED: Line card 10 present
%CHMGR-5-CARDDETECTED: Line card 12 present
%TSM-6-SFM_DISCOVERY: Found SFM 0
%TSM-6-SFM_DISCOVERY: Found SFM 1
%TSM-6-SFM_DISCOVERY: Found SFM 2
%TSM-6-SFM_DISCOVERY: Found SFM 3
%TSM-6-SFM_DISCOVERY: Found SFM 4
%TSM-6-SFM_DISCOVERY: Found SFM 5
%TSM-6-SFM_DISCOVERY: Found SFM 6
%TSM-6-SFM_DISCOVERY: Found SFM 7
%TSM-6-SFM_SWITCHFAB_STATE: Switch Fabric: UP
%TSM-6-SFM_DISCOVERY: Found SFM 8
%TSM-6-SFM_DISCOVERY: Found 9 SFMs
%CHMGR-5-CHECKIN: Checkin from line card 5 (type EX1YB, 1 ports)
%TSM-6-PORT_CONFIG: Port link status for LC 5 => portpipe 0: OK portpipe 1: N/A
%CHMGR-5-LINECARDUP: Line card 5 is up
%CHMGR-5-CHECKIN: Checkin from line card 12 (type S12YC12, 12 ports)
%TSM-6-PORT_CONFIG: Port link status for LC 12 => portpipe 0: OK portpipe 1: N/A
%CHMGR-5-LINECARDUP: Line card 12 is up
%IFMGR-5-CSTATE_UP: changed interface Physical state to up: So 12/8
%IFMGR-5-CSTATE_DN: changed interface Physical state to down: So 12/8
```

To view any changes made, use the `show running-config logging` command in EXEC privilege mode, as shown in the example for [Configuring a UNIX Logging Facility Level](#).

Configuring a UNIX Logging Facility Level

You can save system log messages with a UNIX system logging facility. To configure a UNIX logging facility level, use the following command.

- Specify one of the following parameters.

CONFIGURATION mode

```
logging facility [facility-type]
```

- `auth` (for authorization messages)
- `cron` (for system scheduler messages)
- `daemon` (for system daemons)
- `kern` (for kernel messages)
- `local0` (for local use)
- `local1` (for local use)
- `local2` (for local use)
- `local3` (for local use)
- `local4` (for local use)
- `local5` (for local use)
- `local6` (for local use)
- `local7` (for local use)
- `lpr` (for line printer system messages)
- `mail` (for mail system messages)
- `news` (for USENET news messages)
- `sys9` (system use)
- `sys10` (system use)
- `sys11` (system use)
- `sys12` (system use)
- `sys13` (system use)
- `sys14` (system use)
- `syslog` (for syslog messages)
- `user` (for user programs)
- `uucp` (UNIX to UNIX copy protocol)

Example of the `show running-config logging` Command

To view nondefault settings, use the `show running-config logging` command in EXEC mode.

```
Dell#show running-config logging
!
logging buffered 524288 debugging
service timestamps log datetime msec
```

```
service timestamps debug datetime msec
!
logging trap debugging
logging facility user
logging source-interface Loopback 0
logging 10.10.10.4
Dell#
```

Synchronizing Log Messages

You can configure the system to filter and consolidate the system messages for a specific line by synchronizing the message output.

Only the messages with a severity at or below the set level appear. This feature works on the terminal and console connections available on the system.

1. Enter LINE mode.

CONFIGURATION mode

```
line {console 0 | vty number [end-number]}
```

Configure the following parameters for the virtual terminal lines:

- *number*: the range is from zero (0) to 9.
- *end-number*: the range is from 1 to 8.

You can configure multiple virtual terminals at one time by entering a *number* and an *end-number*.

2. Configure a level and set the maximum number of messages to print.

LINE mode

```
logging synchronous [level severity-level | all] [limit]
```

Configure the following optional parameters:

- *level severity-level*: the range is from 0 to 7. The default is **2**. Use the `all` keyword to include all messages.
- *limit*: the range is from 20 to 300. The default is **20**.

To view the logging synchronous configuration, use the `show config` command in LINE mode.

Enabling Timestamp on Syslog Messages

By default, syslog messages do not include a time/date stamp stating when the error or message was created.

To enable timestamp, use the following command.

- Add timestamp to syslog messages.

CONFIGURATION mode

```
service timestamps [log | debug] [datetime [localtime] [msec] [show-timezone]
| uptime]
```

Specify the following optional parameters:

- `datetime`: You can add the keyword `localtime` to include the `localtime`, `msec`, and `show-timezone`. If you do not add the keyword `localtime`, the time is UTC.
- `uptime`: To view time since last boot.

If you do not specify a parameter, the system configures `uptime`.

To view the configuration, use the `show running-config logging` command in EXEC privilege mode.

To disable time stamping on syslog messages, use the `no service timestamps [log | debug]` command.

File Transfer Services

With the Dell Networking OS, you can configure the system to transfer files over the network using the file transfer protocol (FTP).

One FTP application is copying the system image files over an interface on to the system; however, FTP is not supported on virtual local area network (VLAN) interfaces.

For more information about FTP, refer to RFC 959, *File Transfer Protocol*.

Configuration Task List for File Transfer Services

The configuration tasks for file transfer services are:

- [Enabling the FTP Server](#) (mandatory)
- [Configuring FTP Server Parameters](#) (optional)
- [Configuring FTP Client Parameters](#) (optional)

Enabling the FTP Server

To enable the system as an FTP server, use the following command.

To view FTP configuration, use the `show running-config ftp` command in EXEC privilege mode.

- Enable FTP on the system.
CONFIGURATION mode

```
ftp-server enable
```

Example of Viewing FTP Configuration

```
Dell#show running ftp
!
ftp-server enable
ftp-server username nairobi password 0 zanzibar
Dell#
```

Configuring FTP Server Parameters

After you enable the FTP server on the system, you can configure different parameters. To specify the system logging settings, use the following commands.

- Specify the directory for users using FTP to reach the system.

CONFIGURATION mode

```
ftp-server topdir dir
```

The default is the internal flash directory.

- Specify a user name for all FTP users and configure either a plain text or encrypted password.

CONFIGURATION mode

```
ftp-server username username password [encryption-type] password
```

Configure the following optional and required parameters:

- *username*: enter a text string.
- *encryption-type*: enter 0 for plain text or 7 for encrypted text.
- *password*: enter a text string.



NOTE: You cannot use the `change directory (cd)` command until you have configured `ftp-server topdir`.

To view the FTP configuration, use the `show running-config ftp` command in EXEC privilege mode.

Configuring FTP Client Parameters

To configure FTP client parameters, use the following commands.

- Enter the following keywords and slot/port or number information:
 - For a Loopback interface, enter the keyword `loopback` then a number between 0 and 16383.
 - For a port channel interface, enter the keywords `port-channel` then a number from 1 to 128.
 - For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
 - For a VLAN interface, enter the keyword `vlan` then a number from 1 to 4094.
 - For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

CONFIGURATION mode

```
ip ftp source-interface interface
```

- Configure a password.

CONFIGURATION mode

```
ip ftp password password
```

- Enter a username to use on the FTP client.

CONFIGURATION mode


```
ip ftp username name
```

To view the FTP configuration, use the `show running-config ftp` command in EXEC privilege mode, as shown in the example for [Enabling the FTP Server](#).

Terminal Lines

You can access the system remotely and restrict access to the system by creating user profiles. Terminal lines on the system provide different means of accessing the system. The virtual terminal lines (VTYs) connect you through Telnet to the system.

Denying and Permitting Access to a Terminal Line

Dell Networking recommends applying only standard access control lists (ACLs) to deny and permit access to VTY lines.


- Layer 3 ACLs deny all traffic that is not explicitly permitted, but in the case of VTY lines, an ACL with no rules does not deny traffic.
- You cannot use the `show ip accounting access-list` command to display the contents of an ACL that is applied only to a VTY line.
- When you use the `access-class access-list-name` command without specifying the `ipv4` or `ipv6` attribute, both IPv4 as well as IPv6 rules that are defined in that ACL are applied to the terminal. This is a generic way of configuring access restrictions.
- To be able to filter access exclusively using either IPv4 or IPv6 rules, you must use either the `ipv4` or `ipv6` attribute along with the `access-class access-list-name` command. Depending on the attribute that you specify (`ipv4` or `ipv6`), the ACL processes either IPv4 or IPv6 rules, but not both. Using this configuration, you can set up two different types of access classes with each class processing either IPv4 or IPv6 rules separately.

To apply an IP ACL to a line, Use the following command.

- Apply an ACL to a VTY line.

LINE mode

```
ip access-class access-list [ipv4 | ipv6]
```

 **NOTE:** If you already have configured generic IP ACL on a terminal line, then you cannot further apply IPv4 or IPv6 specific filtering on top of this configuration. Similarly, if you have configured either IPv4 or IPv6 specific filtering on a terminal line, you cannot apply generic IP ACL on top of this configuration. Before applying any of these configurations, you must first undo the existing configuration using the `no access-class access-list-name [ipv4 | ipv6]` command.

Example of an ACL that Permits Terminal Access

To view the configuration, use the `show config` command in LINE mode.

```
Dell(config-std-nacl)#show config
!
ip access-list standard myvtyacl
  seq 5 permit host 10.11.0.1
Dell(config-std-nacl)#line vty 0
Dell(config-line-vty)#show config
```

```
line vty 0
  access-class myvtyacl
```

Dell OS Behavior: Prior to Dell OS version 7.4.2.0, in order to deny access on a VTY line, apply an ACL and accounting, authentication, and authorization (AAA) to the line. Then users are denied access only after they enter a username and password. Beginning in Dell OS version 7.4.2.0, only an ACL is required, and users are denied access before they are prompted for a username and password.

Configuring Login Authentication for Terminal Lines

You can use any combination of up to six authentication methods to authenticate a user on a terminal line.

A combination of authentication methods is called a method list. If the user fails the first authentication method, the system prompts the next method until all methods are exhausted, at which point the connection is terminated. The available authentication methods are:

enable	Prompt for the enable password.
line	Prompt for the password you assigned to the terminal line. Configure a password for the terminal line to which you assign a method list that contains the line authentication method. Configure a password using the <code>password</code> command from LINE mode.
local	Prompt for the system username and password.
none	Do not authenticate the user.
radius	Prompt for a username and password and use a RADIUS server to authenticate.
tacacs+	Prompt for a username and password and use a TACACS+ server to authenticate.

1. Configure an authentication method list. You may use a mnemonic name or use the `default` keyword. The default authentication method for terminal lines is **local** and the default method list is **empty**.

CONFIGURATION mode

```
aaa authentication login {method-list-name | default} [method-1] [method-2]
[method-3] [method-4] [method-5] [method-6]
```

2. Apply the method list from Step 1 to a terminal line.

CONFIGURATION mode

```
login authentication {method-list-name | default}
```

3. If you used the line authentication method in the method list you applied to the terminal line, configure a password for the terminal line.

LINE mode

```
password
```

Example of Terminal Line Authentication

In the following example, VTY lines 0-2 use a single authentication method, line.

```
Dell(conf)#aaa authentication login myvtymethodlist line
Dell(conf)#line vty 0 2
Dell(config-line-vty)#login authentication myvtymethodlist
Dell(config-line-vty)#password myvtypassword
```

```
Dell(config-line-vty)#show config
line vty 0
  password myvtypassword
login authentication myvtymethodlist
line vty 1
  password myvtypassword
login authentication myvtymethodlist
line vty 2
  password myvtypassword
login authentication myvtymethodlist
Dell(config-line-vty)#
```

Setting Time Out of EXEC Privilege Mode

EXEC time-out is a basic security feature that returns the Dell Networking OS to EXEC mode after a period of inactivity on the terminal lines.

To set time out, use the following commands.

- Set the number of minutes and seconds. The default is **10 minutes** on the console and **30 minutes** on VTY. Disable EXEC time out by setting the time-out period to 0.

LINE mode

```
exec-timeout minutes [seconds]
```

- Return to the default time-out values.

LINE mode

```
no exec-timeout
```

Example of Setting the Time Out Period for EXEC Privilege Mode

The following example shows how to set the time-out period and how to view the configuration using the `show config` command from LINE mode.

```
Dell(conf)#line con 0
Dell(config-line-console)#exec-timeout 0
Dell(config-line-console)#show config
line console 0
  exec-timeout 0 0
Dell(config-line-console)#
```

Using Telnet to get to Another Network Device

To telnet to another device, use the following commands.

- Telnet to the stack-unit. You do not need to configure the management port on the stack-unit to be able to telnet to it.

EXEC Privilege mode

```
telnet-peer-stack-unit
```

- Telnet to a device with an IPv4 address.

EXEC Privilege

```
telnet [ip-address]
```

If you do not enter an IP address, the system enters a Telnet dialog that prompts you for one.

Enter an IPv4 address in dotted decimal format (A.B.C.D).

Example of the `telnet` Command for Device Access

```
Dell# telnet 10.11.80.203
Trying 10.11.80.203...
Connected to 10.11.80.203.
Exit character is '^]'.
Login:
Login: admin
Password:
Dell>exit
Dell#telnet 2200:2200:2200:2200:2200::2201
Trying 2200:2200:2200:2200:2200::2201...
Connected to 2200:2200:2200:2200:2200::2201.
Exit character is '^]'.
FreeBSD/i386 (freebsd2.forcel0networks.com) (ttypl)
login: admin
Dell#
```

Lock CONFIGURATION Mode

The systems allows multiple users to make configurations at the same time. You can lock CONFIGURATION mode so that only one user can be in CONFIGURATION mode at any time (Message 2).

You can set two types of locks: auto and manual.

- Set auto-lock using the `configuration mode exclusive auto` command from CONFIGURATION mode. When you set auto-lock, every time a user is in CONFIGURATION mode, all other users are denied access. This means that you can exit to EXEC Privilege mode, and re-enter CONFIGURATION mode without having to set the lock again.
- Set manual lock using the `configure terminal lock` command from CONFIGURATION mode. When you configure a manual lock, which is the default, you must enter this command each time you want to enter CONFIGURATION mode and deny access to others.

Viewing the Configuration Lock Status

If you attempt to enter CONFIGURATION mode when another user has locked it, you may view which user has control of CONFIGURATION mode using the `show configuration lock` command from EXEC Privilege mode.

You can then send any user a message using the `send` command from EXEC Privilege mode.

Alternatively, you can clear any line using the `clear` command from EXEC Privilege mode. If you clear a console session, the user is returned to EXEC mode.


Example of Locking CONFIGURATION Mode for Single-User Access


```
Dell(conf)#configuration mode exclusive auto
BATMAN(conf)#exit
3d23h35m: %RPM0-P:CP %SYS-5-CONFIG_I: Configured from console by console
```

```
Dell#config
! Locks configuration mode exclusively.
Dell(conf)#
```

If another user attempts to enter CONFIGURATION mode while a lock is in place, the following appears on their terminal (message 1): % Error: User "" on line console0 is in exclusive configuration mode.

If *any* user is already in CONFIGURATION mode when while a lock is in place, the following appears on their terminal (message 2): % Error: Can't lock configuration mode exclusively since the following users are currently configuring the system: User "admin" on line vty1 (10.1.1.1).

 **NOTE:** The CONFIGURATION mode lock corresponds to a VTY session, not a user. Therefore, if you configure a lock and then exit CONFIGURATION mode, and another user enters CONFIGURATION mode, when you attempt to re-enter CONFIGURATION mode, you are denied access even though you are the one that configured the lock.

 **NOTE:** If your session times out and you return to EXEC mode, the CONFIGURATION mode lock is unconfigured.

Limit Concurrent Login Sessions

Dell Networking OS enables you to limit the number of concurrent login sessions of users on VTY, Aux, and console lines. You can also clear any of your existing sessions when you reach the maximum permitted number of concurrent sessions.

By default, you can use all 10 VTY lines, one console line, and one Aux line. You can limit the number of available sessions using the `login concurrent-session limit` command and so restrict users to that specific number of sessions. You can optionally configure the system to provide an option to the users to clear any of their existing sessions.

Restrictions for Limiting the Number of Concurrent Sessions

These restrictions apply for limiting the number of concurrent sessions:

- Only the system and security administrators can limit the number of concurrent sessions and enable the clear-line option.
- Users can clear their existing sessions only if the system is configured with the `login concurrent-session clear-line enable` command.

Configuring Concurrent Session Limit

To configure concurrent session limit, follow this procedure:

- Limit the number of concurrent sessions for all users.
CONFIGURATION mode

```
login concurrent-session limit number-of-sessions
```

Example of Configuring Concurrent Session Limit

The following example limits the permitted number of concurrent login sessions to 4.

```
Dell(config)#login concurrent-session limit 4
```

Enabling the System to Clear Existing Sessions

To enable the system to clear existing login sessions, follow this procedure:

- Use the following command.
CONFIGURATION mode

```
login concurrent-session clear-line enable
```

Example of Enabling the System to Clear Existing Sessions

The following example enables you to clear your existing login sessions.

```
Dell(config)#login concurrent-session clear-line enable
```

Example of Clearing Existing Sessions

When you try to login, the following message appears with all your existing concurrent sessions, providing an option to close any one of the existing sessions:

```
$ telnet 10.11.178.14
Trying 10.11.178.14...
Connected to 10.11.178.14.
Escape character is '^]'.
Login: admin
Password:
Current sessions for user admin:
Line          Location
2 vty 0       10.14.1.97
3 vty 1       10.14.1.97
Clear existing session? [line number/Enter to cancel]:
```

When you try to create more than the permitted number of sessions, the following message appears, prompting you to close one of the existing sessions. If you close any of the existing sessions, you are allowed to login. :

```
$ telnet 10.11.178.17
Trying 10.11.178.17...
Connected to 10.11.178.17.
Escape character is '^]'.
Login: admin
Password:

Maximum concurrent sessions for the user reached.
Current VTY sessions for user admin:
Line          Location
2 vty 0       10.14.1.97
3 vty 1       10.14.1.97
4 vty 2       10.14.1.97
5 vty 3       10.14.1.97
Kill existing session? [line number/Enter to cancel]:
```

Track Login Activity

Dell Networking OS enables you to track the login activity of users and view the successful and unsuccessful login events.

When you log in using the console or VTY line, the system displays the last successful login details of the current user and the number of unsuccessful login attempts since your last successful login to the system. The system stores the number of unsuccessful login attempts that have occurred in the last 30 days by default. You can change the default value to any number of days from 1 to 30. By default, login activity tracking is disabled. You can enable it using the `login statistics enable` command from the configuration mode.

Restrictions for Tracking Login Activity

These restrictions apply for tracking login activity:

- Only the system and security administrators can configure login activity tracking and view the login activity details of other users.
- Login statistics is not applicable for login sessions that do not use user names for authentication. For example, the system does not report login activity for a telnet session that prompts only a password.

Configuring Login Activity Tracking

To enable and configure login activity tracking, follow these steps:

1. Enable login activity tracking.

CONFIGURATION mode

```
login statistics enable
```

After enabling login statistics, the system stores the login activity details for the last 30 days.

2. (Optional) Configure the number of days for which the system stores the user login statistics. The range is from 1 to 30.

CONFIGURATION mode

```
login statistics time-period days
```

Example of Configuring Login Activity Tracking

The following example enables login activity tracking. The system stores the login activity details for the last 30 days.

```
Dell(config)#login statistics enable
```

The following example enables login activity tracking and configures the system to store the login activity details for 12 days.

```
Dell(config)#login statistics enable  
Dell(config)#login statistics time-period 12
```

Display Login Statistics

To view the login statistics, use the `show login statistics` command.

Example of the `show login statistics` Command

The `show login statistics` command displays the successful and failed login details of the current user in the last 30 days or the custom defined time period.

```
Dell#show login statistics
```

```
-----  
User: admin  
Last login time: Mon Feb 16 04:40:00 2015  
Last login location: Line vty0 ( 10.14.1.97 )  
Unsuccessful login attempt(s) since the last successful login: 0  
Unsuccessful login attempt(s) in last 30 day(s): 3  
-----
```

Example of the `show login statistics all` command

The `show login statistics all` command displays the successful and failed login details of all users in the last 30 days or the custom defined time period.

```
Dell#show login statistics all
```

```
-----  
User: admin  
Last login time: Mon Feb 16 04:40:00 2015  
Last login location: Line vty0 ( 10.14.1.97 )  
Unsuccessful login attempt(s) since the last successful login: 0  
Unsuccessful login attempt(s) in last 7 day(s): 3  
-----
```

```
-----  
User: secadm  
Last login time: Mon Feb 16 04:45:29 2015  
Last login location: Line vty0 ( 10.14.1.97 )  
Unsuccessful login attempt(s) since the last successful login: 0  
Unsuccessful login attempt(s) in last 7 day(s): 0  
-----
```

Example of the `show login statistics user user-id` command

The `show login statistics user user-id` command displays the successful and failed login details of a specific user in the last 30 days or the custom defined time period.

```
Dell#show login statistics user admin
```

```
-----  
User: admin  
Last login time: Mon Feb 16 04:40:00 2015  
Last login location: Line vty0 ( 10.14.1.97 )  
Unsuccessful login attempt(s) since the last successful login: 0  
Unsuccessful login attempt(s) in last 11 day(s): 3  
-----
```


Recovering from a Forgotten Password

If you configure authentication for the console and you exit out of EXEC mode or your console session times out, you are prompted for a password to re-enter.

Use the following commands if you forget your password.

1. Log onto the system using the console.
2. Power-cycle the chassis by switching off all of the power modules and then switching them back on.
3. Hit any key to abort the boot process. You enter uBoot immediately, as indicated by the => prompt. (during bootup)

hit any key

4. Set the system parameters to ignore the startup configuration file when the system reloads.

uBoot mode

```
setenv stconfigignore true
```

5. To save the changes, use the `saveenv` command.

uBoot mode

```
saveenv
```

6. Reload the system.

uBoot mode

```
reset
```

7. Copy `startup-config.bak` to the running config.

EXEC Privilege mode

```
copy flash://startup-config.bak running-config
```

8. Remove all authentication statements you might have for the console.

LINE mode

```
no authentication login no password
```

9. Save the running-config.

EXEC Privilege mode

```
copy running-config startup-config
```

10. Set the system parameters to use the startup configuration file when the system reloads.

uBoot mode

```
setenv stconfigignore false
```

11. Save the running-config.

EXEC Privilege mode

```
copy running-config startup-config
```

Recovering from a Forgotten Enable Password

Use the following commands if you forget the enable password.

1. Log onto the system using the console.
2. Power-cycle the chassis by switching off all of the power modules and then switching them back on.
3. Hit any key to abort the boot process. You enter uBoot immediately, as indicated by the => prompt. (during bootup)

hit any key

4. Set the system parameters to ignore the enable password when the system reloads.

uBoot mode

```
setenv enablepwdignore true
```

5. Reload the system.

uBoot mode

```
reset
```

6. Configure a new enable password.

CONFIGURATION mode

```
enable {secret | password}
```

7. Save the running-config to the startup-config.

EXEC Privilege mode

```
copy running-config startup-config
```

Recovering from a Failed Start

A system that does not start correctly might be attempting to boot from a corrupted Dell Networking OS image or from a mis-specified location.

In this case, you can restart the system and interrupt the boot process to point the system to another boot location. Use the `setenv` command, as described in the following steps. For details about the `setenv` command, its supporting commands, and other commands that can help recover from a failed start, refer to the *u-Boot* chapter in the *Dell Networking OS Command Line Reference Guide*.

1. Power-cycle the chassis (pull the power cord and reinsert it).
2. Hit any key to abort the boot process. You enter uBoot immediately, the => prompt indicates success. (during bootup)

press any key

3. Assign the new location to the Dell Networking OS image it uses when the system reloads.

uBoot mode

```
setenv [primary_image f10boot location | secondary_image f10boot location |  
default_image f10boot location]
```

4. Assign an IP address to the Management Ethernet interface.

uBoot mode

```
setenv ipaddre address
```

5. Assign an IP address as the default gateway for the system.

uBoot mode

```
setenv gatewayip address
```

6. Reload the system.

uBoot mode


```
reset
```

802.1X

802.1X is a method of port security.

A device connected to a port that is enabled with 802.1X is disallowed from sending or receiving packets on the network until its identity can be verified (through a username and password, for example). This feature is named for its IEEE specification.

802.1X employs extensible authentication protocol (EAP) to transfer a device's credentials to an authentication server (typically RADIUS) using a mandatory intermediary network access device, in this case, a Dell Networking switch. The network access device mediates all communication between the end-user device and the authentication server so that the network remains secure. The network access device uses EAP-over-Ethernet (EAPOL) to communicate with the end-user device and EAP-over-RADIUS to communicate with the server.

 **NOTE:** The Dell Networking operating system supports 802.1X with EAP-MD5, EAP-OTP, EAP-TLS, EAP-TTLS, PEAPv0, PEAPv1, and MS-CHAPv2 with PEAP.

The following figures show how the EAP frames are encapsulated in Ethernet and RADIUS frames.

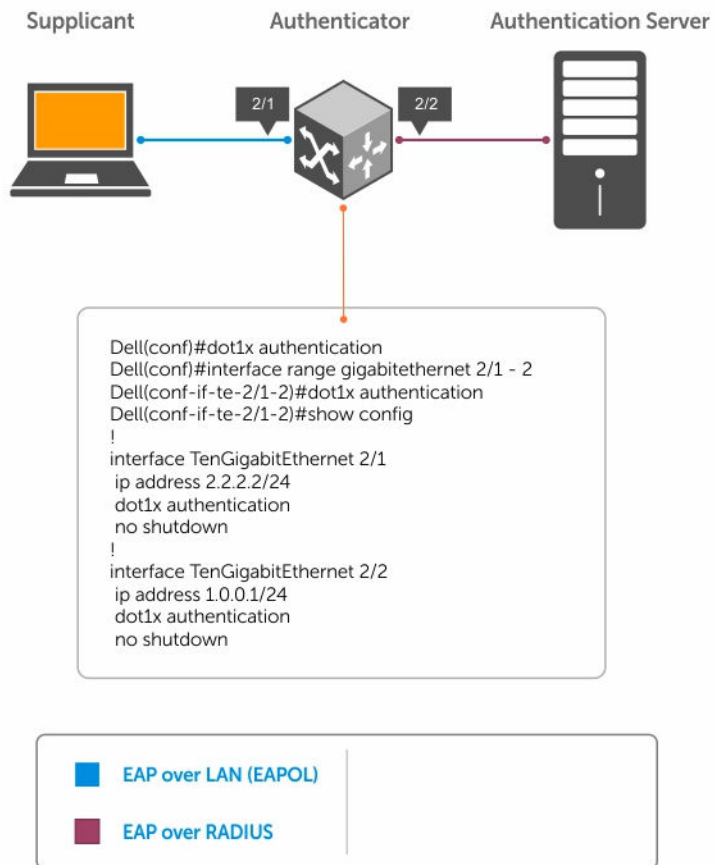


Figure 1. EAP Frames Encapsulated in Ethernet and RADIUS

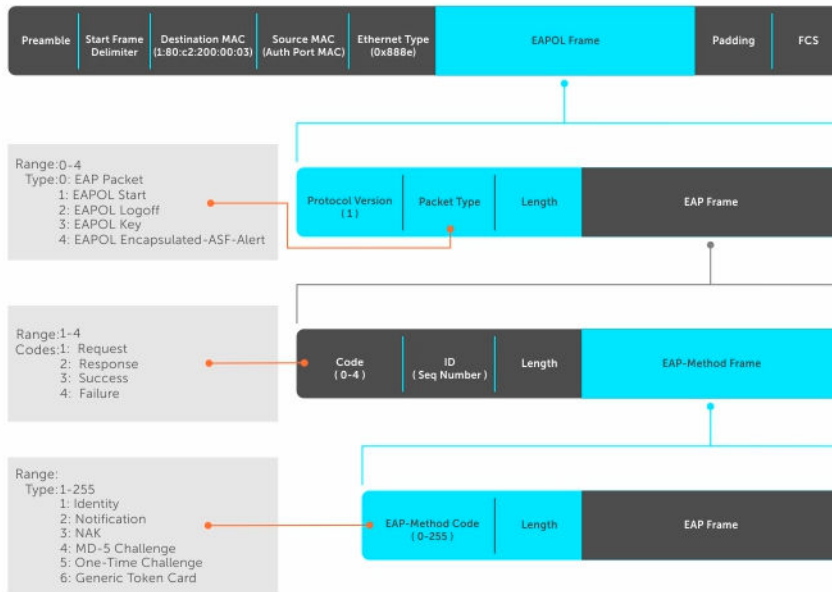


Figure 2. EAP Frames Encapsulated in Ethernet and RADIUS

The authentication process involves three devices:

- The device attempting to access the network is the **supplicant**. The supplicant is not allowed to communicate on the network until the authenticator authorizes the port. It can only communicate with the authenticator in response to 802.1X requests.
- The device with which the supplicant communicates is the **authenticator**. The authenticator is the gate keeper of the network. It translates and forwards requests and responses between the authentication server and the supplicant. The authenticator also changes the status of the port based on the results of the authentication process. The Dell Networking switch is the authenticator.
- The **authentication-server** selects the authentication method, verifies the information the supplicant provides, and grants it network access privileges.

Ports can be in one of two states:

- Ports are in an **unauthorized** state by default. In this state, non-802.1X traffic cannot be forwarded in or out of the port.
- The authenticator changes the port state to **authorized** if the server can authenticate the supplicant. In this state, network traffic can be forwarded normally.

NOTE: The Dell Networking switches place 802.1X-enabled ports in the unauthorized state by default.

The Port-Authentication Process

The authentication process begins when the authenticator senses that a link status has changed from down to up:

1. When the authenticator senses a link state change, it requests that the supplicant identify itself using an EAP Identity Request frame.

- The supplicant responds with its identity in an EAP Response Identity frame.
- The authenticator decapsulates the EAP response from the EAPOL frame, encapsulates it in a RADIUS Access-Request frame and forwards the frame to the authentication server.
- The authentication server replies with an Access-Challenge frame. The Access-Challenge frame requests that the supplicant prove that it is who it claims to be, using a specified method (an EAP-Method). The challenge is translated and forwarded to the supplicant by the authenticator.
- The supplicant can negotiate the authentication method, but if it is acceptable, the supplicant provides the Requested Challenge information in an EAP response, which is translated and forwarded to the authentication server as another Access-Request frame.
- If the identity information provided by the supplicant is valid, the authentication server sends an Access-Accept frame in which network privileges are specified. The authenticator changes the port state to authorized and forwards an EAP Success frame. If the identity information is invalid, the server sends an Access-Reject frame. If the port state remains unauthorized, the authenticator forwards an EAP Failure frame.

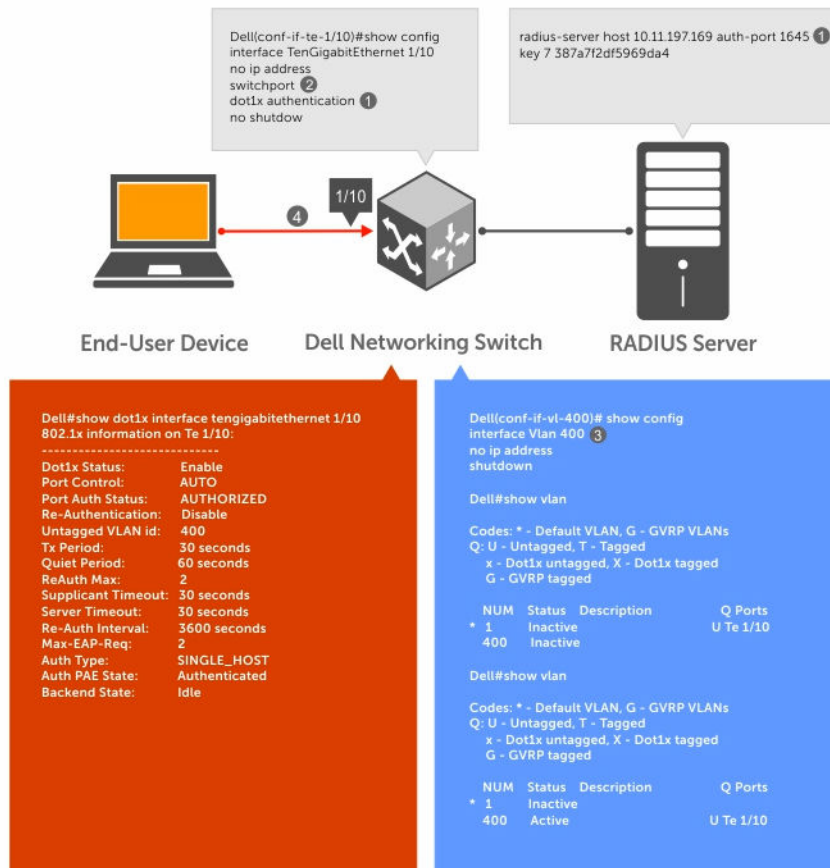


Figure 3. EAP Port-Authentication

EAP over RADIUS

802.1X uses RADIUS to shuttle EAP packets between the authenticator and the authentication server, as defined in RFC 3579.

EAP messages are encapsulated in RADIUS packets as a type of attribute in Type, Length, Value (TLV) format. The Type value for EAP messages is 79.

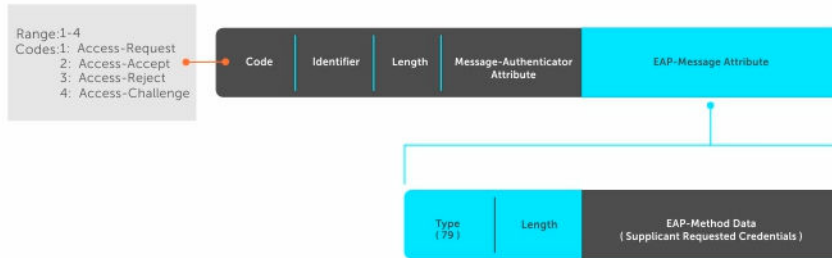


Figure 4. EAP Over RADIUS

RADIUS Attributes for 802.1 Support

Dell Networking systems include the following RADIUS attributes in all 802.1X-triggered Access-Request messages:

- | | |
|---------------------|--|
| Attribute 5 | NAS-Port: the physical port number by which the authenticator is connected to the supplicant. |
| Attribute 31 | Calling-station-id: relays the supplicant MAC address to the authentication server. |
| Attribute 41 | NAS-Port-Type: NAS-port physical port type. 5 indicates Ethernet. |
| Attribute 81 | Tunnel-Private-Group-ID: associate a tunneled session with a particular group of users. |

Configuring 802.1X

Configuring 802.1X on a port is a two-step process.

1. Enable 802.1X globally (refer to [Enabling 802.1X](#)).
2. Enable 802.1X on an interface (refer to [Enabling 802.1X](#)).

Related Configuration Tasks

- [Configuring Request Identity Re-transmissions](#)
- [Forcibly Authorizing or Unauthorizing a Port](#)
- [Re-authenticating a Port](#)
- [Configuring Timeouts](#)

- [Configuring a Guest VLAN](#)
- [Configuring an Authentication-fail VLAN](#)

Important Points to Remember

- The Dell Networking OS supports 802.1X with EAP-MD5, EAP-OTP, EAP-TLS, EAP-TTLS, PEAPv0, PEAPv1, and MS-CHAPv2 with PEAP.
- 802.1X is not supported on port-channels or port-channel members.

Enabling 802.1X

Enable 802.1X globally and at a interface level.

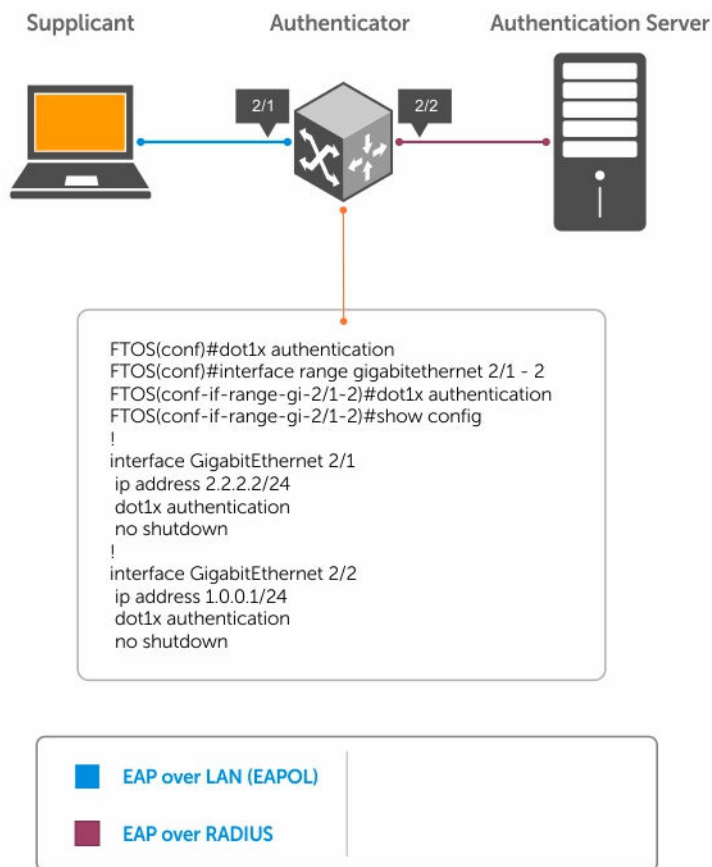


Figure 5. 802.1X Enabled

1. Enable 802.1X globally.
CONFIGURATION mode

```
dot1x authentication
```

2. Enter INTERFACE mode on an interface or a range of interfaces.

```
INTERFACE mode
```

```
interface [range]
```

3. Enable 802.1X on an interface or a range of interfaces.

```
INTERFACE mode
```

```
dot1x authentication
```

Example of Verifying that 802.1X is Enabled Globally

Example of Verifying 802.1X is Enabled on an Interface

Verify that 802.1X is enabled globally and at the interface level using the `show running-config | find dot1x` command from EXEC Privilege mode.

The bold lines show that 802.1X is enabled.

```
Dell#show running-config | find dot1x
dot1x authentication
!
[output omitted]
!
interface GigabitEthernet 2/1
  ip address 2.2.2.2/24
  dot1x authentication
  no shutdown
!
interface GigabitEthernet 2/2
  ip address 1.0.0.1/24
  dot1x authentication
  no shutdown
--More--
```

View 802.1X configuration information for an interface using the `show dot1x interface` command.

The bold lines show that 802.1X is enabled on all ports unauthorized by default.


```
Dell#show dot1x interface TenGigabitEthernet 2/1
```

```
802.1x information on Te 2/1:
-----
Dot1x Status:           Enable
Port Control:             AUTO
Port Auth Status:      UNAUTHORIZED
Re-Authentication:       Disable
Untagged VLAN id:        None
Guest VLAN:              Disable
Guest VLAN id:           NONE
Auth-Fail VLAN:          Disable
Auth-Fail VLAN id:       NONE
Auth-Fail Max-Attempts:  NONE
Mac-Auth-Bypass:         Disable
Mac-Auth-Bypass Only:    Disable
Tx Period:                30 seconds
Quiet Period:             60 seconds
ReAuth Max:               2
Supplicant Timeout:      30 seconds
```

Server Timeout: 30 seconds
Re-Auth Interval: 3600 seconds
Max-EAP-Req: 2
Host Mode: SINGLE_HOST
Auth PAE State: Initialize
Backend State: Initialize

Configuring Request Identity Re-Transmissions

If the authenticator sends a Request Identity frame, but the supplicant does not respond, the authenticator waits 30 seconds and then re-transmits the frame. The amount of time that the authenticator waits before re-transmitting and the maximum number of times that the authenticator re-transmits are configurable.

 **NOTE:** There are several reasons why the supplicant might fail to respond; for example, the supplicant might have been booting when the request arrived or there might be a physical layer problem.

To configure re-transmissions, use the following commands.

- Configure the amount of time that the authenticator waits before re-transmitting an EAP Request Identity frame.

INTERFACE mode

```
dot1x tx-period number
```

The range is from 1 to 65535 (1 year)

The default is **30**.

- Configure a maximum number of times the authenticator re-transmits a Request Identity frame.

INTERFACE mode

```
dot1x max-eap-req number
```


The range is from 1 to 10.

The default is **2**.

The example in [Configuring a Quiet Period after a Failed Authentication](#) shows configuration information for a port for which the authenticator re-transmits an EAP Request Identity frame after 90 seconds and re-transmits a maximum of 10 times.

Configuring a Quiet Period after a Failed Authentication

If the supplicant fails the authentication process, the authenticator sends another Request Identity frame after 30 seconds by default, but you can configure this period.

 **NOTE:** The quiet period (`dot1x quiet-period`) is a transmit interval for after a failed authentication; the Request Identity Re-transmit interval (`dot1x tx-period`) is for an unresponsive supplicant.

To configure a quiet period, use the following command.

- Configure the amount of time that the authenticator waits to re-transmit a Request Identity frame after a failed authentication.

INTERFACE mode

```
dot1x quiet-period seconds
```

The range is from 1 to 65535.

The default is **60 seconds**.

Example of Configuring and Verifying Port Authentication

The following example shows configuration information for a port for which the authenticator re-transmits an EAP Request Identity frame:

- after 90 seconds and a maximum of 10 times for an unresponsive supplicant
- re-transmits an EAP Request Identity frame

The bold lines show the new re-transmit interval, new quiet period, and new maximum re-transmissions.

```
Dell(conf-if-range-Te-0/0)#dot1x tx-period 90
Dell(conf-if-range-Te-0/0)#dot1x max-eap-req 10
Dell(conf-if-range-Te-0/0)#dot1x quiet-period 120
Dell#show dot1x interface TenGigabitEthernet 2/1
802.1x information on Te 2/1:
-----
Dot1x Status:          Enable
Port Control:         AUTO
Port Auth Status:     UNAUTHORIZED
Re-Authentication: Disable
Untagged VLAN id:    None
Tx Period:           90 seconds
Quiet Period:       120 seconds
ReAuth Max:         2
Supplicant Timeout:  30 seconds
Server Timeout:     30 seconds
Re-Auth Interval:   3600 seconds
Max-EAP-Req:       10
Auth Type:          SINGLE_HOST
Auth PAE State:     Initialize
Backend State:      Initialize
```

Forcibly Authorizing or Unauthorizing a Port

IEEE 802.1X requires that a port can be manually placed into any of three states:

- **ForceAuthorized** — an authorized state. A device connected to this port in this state is never subjected to the authentication process, but is allowed to communicate on the network. Placing the port in this state is same as disabling 802.1X on the port.
- **ForceUnauthorized** — an unauthorized state. A device connected to a port in this state is never subjected to the authentication process and is not allowed to communicate on the network. Placing the port in this state is the same as shutting down the port. Any attempt by the supplicant to initiate authentication is ignored.
- **Auto** — an unauthorized state by default. A device connected to this port in this state is subjected to the authentication process. If the process is successful, the port is authorized and the connected device can communicate on the network. All ports are placed in the Auto state by default.

To set the port state, use the following command.

- Place a port in the ForceAuthorized, ForceUnauthorized, or Auto state.

INTERFACE mode

```
dot1x port-control {force-authorized | force-unauthorized | auto}
```

The default state is **auto**.

Example of Placing a Port in Force-Authorized State and Viewing the Configuration

The example shows configuration information for a port that has been force-authorized.

The bold line shows the new port-control state.

```
Dell(conf-if-gi-2/1)#dot1x port-control force-authorized
Dell(conf-if-gi-2/1)#do show dot1x interface gigabitethernet 2/1
```

802.1x information on Gi 2/1:

```
-----
Dot1x Status:          Enable
Port Control:        FORCE AUTHORIZED
Port Auth Status:     UNAUTHORIZED
Re-Authentication:    Disable
Untagged VLAN id:     None
Tx Period:            90 seconds
Quiet Period:         120 seconds
ReAuth Max:           2
Supplicant Timeout:   30 seconds
Server Timeout:       30 seconds
Re-Auth Interval:     3600 seconds
Max-EAP-Req:          10
Auth Type:             SINGLE_HOST

Auth PAE State:       Initialize
Backend State:        Initialize
Auth PAE State:       Initialize
Backend State:        Initialize
```

Re-Authenticating a Port

You can configure the authenticator for periodic re-authentication.

After the supplicant has been authenticated, and the port has been authorized, you can configure the authenticator to re-authenticate the supplicant periodically. If you enable re-authentication, the supplicant is required to re-authenticate every 3600 seconds, but you can configure this interval. You can configure a maximum number of re-authentications as well.

To configure re-authentication time settings, use the following commands.

- Configure the authenticator to periodically re-authenticate the supplicant.

INTERFACE mode

```
dot1x reauthentication [interval] seconds
```

The range is from 1 to 65535.

The default is **3600**.

- Configure the maximum number of times that the supplicant can be re-authenticated.

INTERFACE mode

```
dot1x reauth-max number
```

The range is from 1 to 10.

The default is **2**.

Example of Re-Authenticating a Port and Verifying the Configuration

The bold lines show that re-authentication is enabled and the new maximum and re-authentication time period.

```
Dell(conf-if-gi-2/1)#dot1x reauthentication interval 7200
Dell(conf-if-gi-2/1)#dot1x reauth-max 10
Dell(conf-if-gi-2/1)#do show dot1x interface gigabitethernet 2/1
```

802.1x information on Gi 2/1:

```
-----
Dot1x Status:          Enable
Port Control:         FORCE AUTHORIZED
Port Auth Status:UNAUTHORIZED
Re-Authentication:    Enable
Untagged VLAN id:     None
Tx Period:            90 seconds
Quiet Period:         120 seconds
ReAuth Max:          10
Supplicant Timeout:   30 seconds
Server Timeout:       30 seconds
Re-Auth Interval:7200 seconds
Max-EAP-Req:          10
Auth Type:             SINGLE HOST
Auth PAE State:        Initialize
Backend State:         Initialize
Auth PAE State:        Initialize
```

Configuring Timeouts

If the supplicant or the authentication server is unresponsive, the authenticator terminates the authentication process after 30 seconds by default. You can configure the amount of time the authenticator waits for a response.

To terminate the authentication process, use the following commands.

- Terminate the authentication process due to an unresponsive supplicant.

INTERFACE mode

```
dot1x supplicant-timeout seconds
```

The range is from 1 to 300.

The default is **30**.

- Terminate the authentication process due to an unresponsive authentication server.

INTERFACE mode

```
dot1x server-timeout seconds
```

The range is from 1 to 300.

The default is **30**.

Example of Viewing Configured Server Timeouts

The example shows configuration information for a port for which the authenticator terminates the authentication process for an unresponsive supplicant or server after 15 seconds.

The bold lines show the new supplicant and server timeouts.

```
Dell(conf-if-gi-2/1)#dot1x port-control force-authorized
Dell(conf-if-gi-2/1)#do show dot1x interface gigabitethernet 2/1

802.1x information on Gi 2/1:
-----
Dot1x Status:          Enable
Port Control:          FORCE_AUTHORIZED
Port Auth Status:      UNAUTHORIZED
Re-Authentication:     Disable
Untagged VLAN id:      None
Guest VLAN:            Disable
Guest VLAN id:         NONE
Auth-Fail VLAN:        Disable
Auth-Fail VLAN id:     NONE
Auth-Fail Max-Attempts: NONE
Tx Period:             90 seconds
Quiet Period:          120 seconds
ReAuth Max:            10
Supplicant Timeout: 15 seconds
Server Timeout:        15 seconds
Re-Auth Interval:      7200 seconds
Max-EAP-Req:           10
Auth Type:             SINGLE_HOST
Auth PAE State:        Initialize
Backend State:         Initialize
```

Enter the tasks the user should do after finishing this task (optional).

Configuring Dynamic VLAN Assignment with Port Authentication

The system supports dynamic VLAN assignment when using 802.1X.

The basis for VLAN assignment is RADIUS attribute 81, Tunnel-Private-Group-ID. Dynamic VLAN assignment uses the standard dot1x procedure:

1. The host sends a dot1x packet to the Dell Networking system
2. The system forwards a RADIUS REQUEST packet containing the host MAC address and ingress port number
3. The RADIUS server authenticates the request and returns a RADIUS ACCEPT message with the VLAN assignment using Tunnel-Private-Group-ID

The illustration shows the configuration on the Dell Networking system before connecting the end user device in black and blue text, and after connecting the device in red text. The blue text corresponds to the preceding numbered steps on dynamic VLAN assignment with 802.1X.

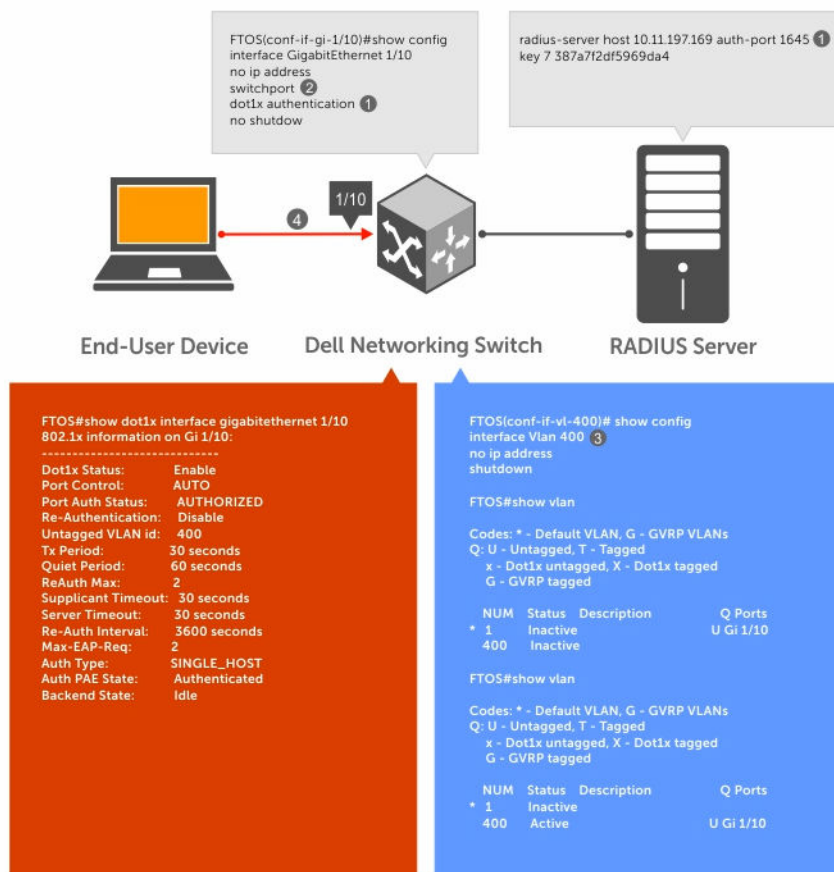



Figure 6. Dynamic VLAN Assignment

1. Configure 802.1x globally (refer to [Enabling 802.1X](#)) along with relevant RADIUS server configurations (refer to the illustration in [Dynamic VLAN Assignment with Port Authentication](#)).
2. Make the interface a switchport so that it can be assigned to a VLAN.
3. Create the VLAN to which the interface will be assigned.
4. Connect the supplicant to the port configured for 802.1X.
5. Verify that the port has been authorized and placed in the desired VLAN (refer to the illustration in [Dynamic VLAN Assignment with Port Authentication](#)).

Guest and Authentication-Fail VLANs

Typically, the authenticator (the Dell Networking system) denies the supplicant access to the network until the supplicant is authenticated. If the supplicant is authenticated, the authenticator enables the port and places it in either the VLAN for which the port is configured or the VLAN that the authentication server indicates in the authentication data.

 **NOTE:** Ports cannot be dynamically assigned to the default VLAN.


If the supplicant fails authentication, the authenticator typically does not enable the port. In some cases this behavior is not appropriate. External users of an enterprise network, for example, might not be able to be authenticated, but still need access to the network. Also, some dumb-terminals, such as network printers, do not have 802.1X capability and therefore cannot authenticate themselves. To be able to connect such devices, they must be allowed access the network without compromising network security.

The Guest VLAN 802.1X extension addresses this limitation with regard to non-802.1X capable devices and the Authentication-fail VLAN 802.1X extension addresses this limitation with regard to external users.

- If the supplicant fails authentication a specified number of times, the authenticator places the port in the Authentication-fail VLAN.
- If a port is already forwarding on the Guest VLAN when 802.1X is enabled, the port is moved out of the Guest VLAN and the authentication process begins.

Configuring a Guest VLAN

If the supplicant does not respond within a determined amount of time ($[reauth-max + 1] * tx-period$), the system assumes that the host does not have 802.1X capability and the port is placed in the Guest VLAN.

 **NOTE:** For more information about configuring timeouts, refer to [Configuring Timeouts](#).


Configure a port to be placed in the Guest VLAN after failing to respond within the timeout period using the `dot1x guest-vlan` command from INTERFACE mode. View your configuration using the `show config` command from INTERFACE mode or using the `show dot1x interface` command from EXEC Privilege mode.

Example of Viewing Guest VLAN Configuration

```
Dell(conf-if-gi-1/2)#dot1x guest-vlan 200
Dell(conf-if-gi-1/2)#show config
!
interface GigabitEthernet 1/2
  switchport
  dot1x guest-vlan 200
  no shutdown
Dell(conf-if-gi-1/2)#
```

Configuring an Authentication-Fail VLAN

If the supplicant fails authentication, the authenticator re-attempts to authenticate after a specified amount of time.

 **NOTE:** For more information about authenticator re-attempts, refer to [Configuring a Quiet Period after a Failed Authentication](#)

You can configure the maximum number of times the authenticator re-attempts authentication after a failure (**3** by default), after which the port is placed in the Authentication-fail VLAN.

Configure a port to be placed in the VLAN after failing the authentication process as specified number of times using the `dot1x auth-fail-vlan` command from INTERFACE mode. Configure the maximum number of authentication attempts by the authenticator using the keyword `max-attempts` with this command.

Example of Configuring Maximum Authentication Attempts

Example of Viewing Configured Authentication

```
Dell(conf-if-gi-1/2)#dot1x auth-fail-vlan 100 max-attempts 5
Dell(conf-if-gi-1/2)#show config
!
interface GigabitEthernet 1/2
  switchport
  dot1x guest-vlan 200
  dot1x auth-fail-vlan 100 max-attempts 5
  no shutdown
Dell(conf-if-gi-1/2)#
```

View your configuration using the `show config` command from INTERFACE mode, as shown in the example in [Configuring a Guest VLAN](#) or using the `show dot1x interface` command from EXEC Privilege mode.

```
Dell(conf-if-gi-2/1)#dot1x port-control force-authorized
Dell(conf-if-gi-2/1)#do show dot1x interface gigabitethernet 2/1
```

```
802.1x information on Gi 2/1:
-----
Dot1x Status:          Enable
Port Control:         FORCE_AUTHORIZED
Port Auth Status:     UNAUTHORIZED
Re-Authentication:    Disable
Untagged VLAN id:     None
Guest VLAN:          Enable
Guest VLAN id:      200
Auth-Fail VLAN:     Enable
Auth-Fail VLAN id:  100
Auth-Fail Max-Attempts: 5
Tx Period:            90 seconds
Quiet Period:         120 seconds
ReAuth Max:           10
Supplicant Timeout:   15 seconds
Server Timeout:       15 seconds
Re-Auth Interval:     7200 seconds
Max-EAP-Req:          10
Auth Type:            SINGLE_HOST
Auth PAE State:        Initialize
Backend State:         Initialize
```

Access Control List (ACL) VLAN Groups and Content Addressable Memory (CAM)

This chapter describes the access control list (ACL) VLAN group and content addressable memory (CAM) enhancements.

Optimizing CAM Utilization During the Attachment of ACLs to VLANs

You can enable and configure the ACL CAM optimization functionality to minimize the number of entries in CAM while ACLs are applied on a VLAN or a set of VLANs, and also while ACLs are applied on a set of ports. This capability enables the effective usage of the CAM space when Layer 3 ACLs are applied to a set of VLANs and when Layer 2 or Layer 3 ACLs are applied on a set of ports.

In releases of Dell Networking OS that do not support the CAM optimization functionality, when an ACL is applied on a VLAN, the ACL rules are configured with the rule-specific parameters and the VLAN as additional attributes in the ACL region. When the ACL is applied on multiple VLAN interfaces, the consumption of the CAM space increases proportionally. For example, when an ACL with 'n' number of rules is applied on 'm' number of VLAN interfaces, a total of n*m entries are configured in the CAM region that is allocated for ACLs. Similarly, when an L2 or L3 ACL is applied on a set of ports, a large portion of the CAM space gets used because a port is saved as a parameter in CAM.

To avoid excessive consumption of the CAM space, configure ACL VLAN groups, which combine all the VLANs that are applied with the same ACL, into a single group. A class identifier (Class ID) is assigned for each of the ACLs attached to the VLAN and this Class ID is used as an identifier or locator in the CAM space instead of the VLAN ID. This method of processing reduces the number of entries in the CAM area significantly and saves memory space by using the class ID as a filtering criterion in CAM instead of the VLAN ID.

You can create an ACL VLAN group and attach the ACL with the VLAN members. The optimization is applicable only when you create an ACL VLAN group. If you apply an ACL separately on the VLAN interface, each ACL has a mapping with the VLAN and increased CAM space utilization occurs. Attaching an ACL individually to VLAN interfaces is similar to the behavior of ACL-VLAN mapping storage in CAM prior to the implementation of the ACL VLAN group functionality.

The ACL manager application on router processor (RP1) contains all the state information about all the ACL VLAN groups that are present. The ACL handler on control processor (CP) and the ACL agent on line cards do not contain any stateful information about the group. The ACL manager application performs the validation after you enter the `acl-vlan-group` command. If the command is valid, it is processed and sent to the agent, if required. If a configuration error is found or if the maximum limit has exceeded

for the ACL VLAN groups present on the system, an appropriate error message is displayed. The ACL manager application verifies the following parameters when you enter the `acl-vlan-group` command:

- Whether the CAM profile is set in VFP
- Whether the maximum number of groups in the system has exceeded
- Whether the maximum number of VLAN numbers permitted per ACL group has exceeded
- When a VLAN member that is being added is already a part of another ACL group

After these verification steps are performed, the ACL manager considers the command as valid and sends the information to the ACL agent on the line card. The ACL manager notifies the ACL agent in the following cases:

- A VLAN member is added or removed from a group, and previously associated VLANs exist in the group.
- The egress ACL is applied or removed from the group and the group contains VLAN members. VLAN members are added or deleted from a VLAN, which itself is a group member.
- A line card returns to the active state after going down, and this line card contains a VLAN that is a member of an ACL group.
- The ACL VLAN group is deleted and it contains VLAN members.

The ACL manager does not notify the ACL agent in the following cases:

- The ACL VLAN group is created.
- The ACL VLAN group is deleted and it does not contain any VLAN members.
- The ACL is applied or removed from a group, and the ACL group does not contain a VLAN member.
- The description of the ACL group is added or removed.

Guidelines for Configuring ACL VLAN groups

Keep the following points in mind when you configure ACL VLAN groups:

- The interfaces, to which the ACL VLAN group is applied, function as restricted interfaces. The ACL VLAN group name is used to identify the group of VLANs that is used to perform hierarchical filtering.
- You can add only one ACL to an interface at a time.
- When you attach an ACL VLAN group to the same interface, a validation is performed to determine whether an ACL is applied directly to an interface. If you previously applied an ACL separately to the interface, an error occurs when you attempt to attach an ACL VLAN group to the same interface.
- The maximum number of members in an ACL VLAN group is determined by the type of switch and its hardware capabilities. This scaling limit depends on the number of slices that are allocated for ACL CAM optimization. If one slice is allocated, the maximum number of VLAN members is 256 for all ACL VLAN groups. If two slices are allocated, the maximum number of VLAN members is 512 for all ACL VLAN groups.
- The maximum number of VLAN groups that you can configure also depends on the hardware specifications of the switch. Each VLAN group is mapped to a unique ID in the hardware. The maximum number of ACL VLAN groups supported is 31. Only a maximum of two components (iSCSI counters, Open Flow, ACL optimization) can be allocated virtual flow processing slices at a time.

- The maximum number of VLANs that you can configure as a member of ACL VLAN groups is limited to 512 on the MXL switch if two slices are allocated. If only one virtual flow processing slice is allocated, the maximum number of VLANs that you can configure as a member of an ACL VLAN group is 256 for the MXL switch.
- Port ACL optimization is applicable only for ACLs that are applied without the VLAN range.
- You cannot view the statistical details of ACL rules per VLAN and per interface if you enable the ACL VLAN group capability. You can view the counters per ACL only using the `show ip accounting access list` command.
- Within a port, you can apply Layer 2 ACLs on a VLAN or a set of VLANs. In this case, CAM optimization is not applied.
- To enable optimization of CAM space for Layer 2 or Layer 3 ACLs that are applied to ports, the port number is removed as a qualifier for ACL application on ports, and port bits are used. When you apply the same ACL to a set of ports, the port bitmap is set when the ACL flow processor (FP) entry is added. When you remove the ACL from a port, the port bitmap is removed.
- If you do not attach an ACL to any of the ports, the FP entries are deleted. Similarly, when the same ACL is applied on a set of ports, only one set of entries is installed in the FP, thereby effectively saving CAM space. The optimization is enabled only if you specify the optimized option with the `ip access-group` command. This option is not valid for VLAN and LAG interfaces.

Configuring ACL VLAN Groups and Configuring FP Blocks for VLAN Parameters

This section describes how to optimize the utilization of CAM blocks by configuring ACL VLAN groups that you can attach to VLAN interfaces and also how to configure FP blocks for different VLAN operations.

Configuring ACL VLAN Groups

You can create an ACL VLAN group and attach the ACL with the VLAN members. The optimization is applicable only when you create an ACL VLAN group. If you apply an ACL separately on the VLAN interface, each ACL has a mapping with the VLAN and increases the CAM space utilization. Attaching an ACL individually to VLAN interfaces is similar to the behavior of ACL-VLAN mapping storage in CAM prior to the implementation of the ACL VLAN group functionality.

1. Create an ACL VLAN group

CONFIGURATION mode

```
acl-vlan-group {group name}
```

You can have up to eight different ACL VLAN groups at any given time.

2. Add a description to the ACL VLAN group.

CONFIGURATION (conf-acl-vl-grp) mode

```
description description
```

3. Apply an egress IP ACL to the ACL VLAN group.

CONFIGURATION (conf-acl-vl-grp) mode

- ```
ip access-group {group name} out implicit-permit
```
4. Add VLAN member(s) to an ACL VLAN group.  
CONFIGURATION (conf-acl-vl-grp) mode  
  
member vlan {VLAN-range}
  5. Display all the ACL VLAN groups or display a specific ACL VLAN group, identified by name.  
CONFIGURATION (conf-acl-vl-grp) mode  
  
show acl-vlan-group {group name | detail}  
Dell#show acl-vlan-group detail  
  
Group Name :  
 TestGroupSeventeenTwenty  
Egress IP Acl :  
 SpecialAccessOnlyExpertsAllowed  
Vlan Members :  
 100,200,300  
  
Group Name :  
 CustomerNumberIdentificationEleven  
Egress IP Acl :  
 AnyEmployeeCustomerElevenGrantedAccess  
Vlan Members :  
 2-10,99  
  
Group Name :  
 HostGroup  
Egress IP Acl :  
 Group5  
Vlan Members :  
 1,1000  
Dell#

## Configuring FP Blocks for VLAN Parameters

Use the `cam-acl-vlan` command to allocate the number of FP blocks for the various VLAN processes on the system. You can use the **no** version of this command to reset the number of FP blocks to default. By default, 0 groups are allocated for the ACL in VCAP. ACL VLAN groups or CAM optimization is not enabled by default, and you need to allocate the slices for CAM optimization.

1. Allocate the number of FP blocks for VLAN Open Flow operations.  
CONFIGURATION mode  
  
cam-acl-vlan vlanopenflow <0-2>
2. Allocate the number of FP blocks for VLAN iSCSI counters.  
CONFIGURATION mode  
  
cam-acl-vlan vlaniscsi <0-2>
3. Allocate the number of FP blocks for ACL VLAN optimization feature.  
CONFIGURATION mode  
  
cam-acl-vlan vlna clopt <0-2>
4. View the number of flow processor (FP) blocks that is allocated for the different VLAN services.

EXEC Privilege mode

Dell#show cam-usage switch

| Linecard | Portpipe | CAM Partition | Total CAM | Used CAM | Available CAM |
|----------|----------|---------------|-----------|----------|---------------|
| 11       | 0        | IN-L2 ACL     | 7152      | 0        |               |
| 7152     |          | IN-L2 FIB     | 32768     | 1081     |               |
| 31687    |          | OUT-L2 ACL    | 0         | 0        |               |
| 0        |          |               |           |          |               |
| 11       | 1        | IN-L2 ACL     | 7152      | 0        |               |
| 7152     |          | IN-L2 FIB     | 32768     | 1081     |               |
| 31687    |          | OUT-L2 ACL    | 0         | 0        | 0             |

## Viewing CAM Usage

View the amount of CAM space available, used, and remaining in each partition (including IPv4Flow and Layer 2 ACL sub- partitions) using the `show cam-usage` command in EXEC Privilege mode

Display Layer 2, Layer 3, ACL, or all CAM usage statistics.

EXCE Privilege mode

```
show cam usage [acl | router | switch]
```

The following sample output shows the consumption of CAM blocks for Layer 2 and Layer 3 ACLs, in addition to other processes that use CAM space:

```
Dell#show cam-usage
```

| Linecard | Portpipe | CAM Partition  | Total CAM | Used CAM | Available CAM |
|----------|----------|----------------|-----------|----------|---------------|
| 1        | 0        | IN-L2 ACL      | 1008      | 320      | 688           |
|          |          | IN-L2 FIB      | 32768     | 1132     | 31636         |
|          |          | IN-L3 ACL      | 12288     | 2        | 12286         |
|          |          | IN-L3 FIB      | 262141    | 14       | 262127        |
|          |          | IN-L3-SysFlow  | 2878      | 45       | 2833          |
|          |          | IN-L3-TrcList  | 1024      | 0        | 1024          |
|          |          | IN-L3-McastFib | 9215      | 0        | 9215          |
|          |          | IN-L3-Qos      | 8192      | 0        | 8192          |
|          |          | IN-L3-PBR      | 1024      | 0        | 1024          |
|          |          | IN-V6 ACL      | 0         | 0        | 0             |
|          |          | IN-V6 FIB      | 0         | 0        | 0             |
|          |          | IN-V6-SysFlow  | 0         | 0        | 0             |
|          |          | IN-V6-McastFib | 0         | 0        | 0             |
|          |          | OUT-L2 ACL     | 1024      | 0        | 1024          |
|          |          | OUT-L3 ACL     | 1024      | 0        | 1024          |
|          |          | OUT-V6 ACL     | 0         | 0        | 0             |
| 1        | 1        | IN-L2 ACL      | 320       | 0        | 320           |
|          |          | IN-L2 FIB      | 32768     | 1136     | 31632         |
|          |          | IN-L3 ACL      | 12288     | 2        | 12286         |
|          |          | IN-L3 FIB      | 262141    | 14       | 262127        |
|          |          | IN-L3-SysFlow  | 2878      | 44       | 2834          |

--More--

The following sample output displays the CAM space utilization when Layer 2 and Layer 3 ACLs are configured:

```
Dell#show cam-usage acl
Linecard|Portpipe| CAM Partition | Total CAM | Used CAM | Available CAM
=====|=====|=====|=====|=====|=====
 11 | 0 | IN-L2 ACL | 1008 | 0 | 1008
 | | IN-L3 ACL | 12288 | 2 | 12286
 | | OUT-L2 ACL | 1024 | 2 | 1022
 | | OUT-L3 ACL | 1024 | 0 | 1024
```

The following sample output displays the CAM space utilization for Layer 2 ACLs:

```
Dell#show cam-usage switch
Linecard|Portpipe| CAM Partition | Total CAM | Used CAM | Available CAM
=====|=====|=====|=====|=====|=====
 11 | 0 | IN-L2 ACL | 7152 | 0 | 7152
 | | IN-L2 FIB | 32768 | 1081 | 31687
 | | OUT-L2 ACL | 0 | 0 | 0
 11 | 1 | IN-L2 ACL | 7152 | 0 | 7152
 | | IN-L2 FIB | 32768 | 1081 | 31687
 | | OUT-L2 ACL | 0 | 0 | 0
```

The following sample output displays the CAM space utilization for Layer 3 ACLs:

```
Dell#show cam-usage router
Linecard|Portpipe| CAM Partition | Total CAM | Used CAM | Available CAM
=====|=====|=====|=====|=====|=====
 11 | 0 | IN-L3 ACL | 8192 | 3 | 8189
 | | IN-L3 FIB | 196607 | 1 | 196606
 | | IN-L3-SysFlow | 2878 | 0 | 2878
 | | IN-L3-TrcList | 1024 | 0 | 1024
 | | IN-L3-McastFib | 9215 | 0 | 9215
 | | IN-L3-Qos | 8192 | 0 | 8192
 | | IN-L3-PBR | 1024 | 0 | 1024
 | | OUT-L3 ACL | 16384 | 0 | 16384
 11 | 1 | IN-L3 ACL | 8192 | 3 | 8189
 | | IN-L3 FIB | 196607 | 1 | 196606
 | | IN-L3-SysFlow | 2878 | 0 | 2878
 | | IN-L3-TrcList | 1024 | 0 | 1024
 | | IN-L3-McastFib | 9215 | 0 | 9215
 | | IN-L3-Qos | 8192 | 0 | 8192
 | | IN-L3-PBR | 1024 | 0 | 1024
 | | OUT-L3 ACL | 16384 | 0 | 16384
```

## Allocating FP Blocks for VLAN Processes

The VLAN ContentAware Processor (VCAP) application is a preingress CAP that modifies the VLAN settings before packets are forwarded. To support the ACL CAM optimization functionality, the CAM carving feature is enhanced. A total of four VACP groups are present, of which two are for fixed groups and the other two are for dynamic groups. Out of the total of two dynamic groups, you can allocate zero, one, or two FP blocks to iSCSI Counters, OpenFlow and ACL Optimization.



You can configure only two of these features at a time.

- To allocate the number of FP blocks for VLAN open flow operations, use the `cam-acl-vlan vlanopenflow <0-2>` command.
- To allocate the number of FP blocks for VLAN iSCSI counters, use the `cam-acl-vlan vlaniscsi <0-2>` command.
- To allocate the number of FP blocks for ACL VLAN optimization feature, use the `cam-acl-vlan vlnaopt <0-2>` command.

To reset the number of FP blocks to the default, use the **no** version of these commands. By default, zero groups are allocated for the ACL in VCAP. ACL VLAN groups or CAM optimization is not enabled by default, and you need to allocate the slices for CAM optimization.

To display the number of FP blocks that is allocated for the different VLAN services, you can use the `show cam-acl-vlan` command. After CAM configuration for ACL VLAN groups is performed, reboot the system to enable the settings to be stored in nonvolatile storage. During the initialization of CAM, the chassis manager reads the NVRAM and allocates the dynamic VCAP regions.

## Access Control Lists (ACLs)

This chapter describes access control lists (ACLs), prefix lists, and route-maps.

At their simplest, ACLs, prefix lists, and route-maps permit or deny traffic based on MAC and/or IP addresses. This chapter describes implementing IP ACLs, IP prefix lists and route-maps. For MAC ACLs, refer to [Layer 2](#).

An ACL is essentially a filter containing some criteria to match (examine IP, transmission control protocol [TCP], or user datagram protocol [UDP] packets) and an action to take (permit or deny). ACLs are processed in sequence so that if a packet does not match the criterion in the first filter, the second filter (if configured) is applied. When a packet matches a filter, the switch drops or forwards the packet based on the filter's specified action. If the packet does not match any of the filters in the ACL, the packet is dropped (implicit deny).

The number of ACLs supported on a system depends on your content addressable memory (CAM) size. For more information, refer to the [Content Addressable Memory \(CAM\)](#) chapter.

## IP Access Control Lists (ACLs)

In Dell Networking switch/routers, you can create two different types of IP ACLs: standard or extended.

A standard ACL filters packets based on the source IP packet. An extended ACL filters traffic based on the following criteria:

- IP protocol number
- Source IP address
- Destination IP address
- Source TCP port number
- Destination TCP port number
- Source UDP port number
- Destination UDP port number


For more information about ACL options, refer to the *Dell Networking OS Command Reference Guide*.

For extended ACL, TCP, and UDP filters, you can match criteria on specific or ranges of TCP or UDP ports. For extended ACL TCP filters, you can also match criteria on established TCP sessions.

When creating an access list, the sequence of the filters is important. You have a choice of assigning sequence numbers to the filters as you enter them, or the Dell Networking operating system assigns numbers in the order the filters are created. The sequence numbers are listed in the display output of the `show config` and `show ip accounting access-list` commands.

Ingress and egress hot lock ACLs allow you to append or delete new rules into an existing ACL (already written into CAM) without disrupting traffic flow. Existing entries in the CAM are shuffled to

accommodate the new entries. Hot lock ACLs are enabled by default and support both standard and extended ACLs.

 **NOTE:** Hot lock ACLs are supported for Ingress ACLs only.


## Implementing ACL on the Dell Networking OS

You can assign one IP ACL per interface with the Dell Networking OS. If you do not assign an IP ACL to an interface, it is not used by the software in any other capacity.

The number of entries allowed per ACL is hardware-dependent. For detailed specification on entries allowed per ACL, refer to your line card documentation.

If you enable counters on IP ACL rules that are already configured, those counters are reset when a new rule is inserted or prepended. If a rule is appended, the existing counters are not affected. This is applicable to the following features:

- L2 Ingress Access list
- L2 Egress Access list
- L3 Ingress Access list
- L3 Egress Access list

 **NOTE:** IP ACLs are supported over VLANs in the Dell Networking OS version 6.2.1.1 and higher.

## ACLs and VLANs

There are some differences when assigning ACLs to a VLAN rather than a physical port.

For example, when using a single port-pipe, if you apply an ACL to a VLAN, one copy of the ACL entries is installed in the ACL CAM on the port-pipe. The entry looks for the incoming VLAN in the packet. Whereas if you apply an ACL on individual ports of a VLAN, separate copies of the ACL entries are installed for each port belonging to a port-pipe.

## ACL Optimization

If an access list contains duplicate entries, the system deletes one entry to conserve CAM space.

Standard and extended ACLs take up the same amount of CAM space. A single ACL rule uses two CAM entries whether it is identified as a standard or extended ACL.

## Determine the Order in which ACLs are Used to Classify Traffic

When you link class-maps to queues using the `service-queue` command, the system matches the class-maps according to queue priority (queue numbers closer to 0 have lower priorities).

As shown in the following example, class-map `cmap2` is matched against ingress packets before `cmap1`.

ACLs *acl1* and *acl2* have overlapping rules because the address range 20.1.1.0/24 is within 20.0.0.0/8. Therefore (without the keyword *order*), packets within the range 20.1.1.0/24 match positive against *cmap1* and are buffered in queue 7, though you intended for these packets to match positive against *cmap2* and be buffered in queue 4.

In cases such as these, where class-maps with overlapping ACL rules are applied to different queues, use the *order* keyword to specify the order in which you want to apply ACL rules. The order can range from 0 to 254. The Dell Networking OS writes to the CAM ACL rules with lower-order numbers (order numbers closer to 0) before rules with higher-order numbers so that packets are matched as you intended. By default, all ACL rules have an order of **255**.

## Example of the *order* Keyword to Determine ACL Sequence

```
Dell(conf)#ip access-list standard acl1
Dell(config-std-nacl)#permit 20.0.0.0/8
Dell(config-std-nacl)#exit
Dell(conf)#ip access-list standard acl2
Dell(config-std-nacl)#permit 20.1.1.0/24 order 0
Dell(config-std-nacl)#exit
Dell(conf)#class-map match-all cmap1
Dell(conf-class-map)#match ip access-group acl1
Dell(conf-class-map)#exit
Dell(conf)#class-map match-all cmap2
Dell(conf-class-map)#match ip access-group acl2
Dell(conf-class-map)#exit
Dell(conf)#policy-map-input pmap
Dell(conf-policy-map-in)#service-queue 7 class-map cmap1
Dell(conf-policy-map-in)#service-queue 4 class-map cmap2
Dell(conf-policy-map-in)#exit
Dell(conf)#interface gig 1/0
Dell(conf-if-gi-1/0)#service-policy input pmap
```

## IP Fragment Handling

The Dell Networking OS supports a configurable option to explicitly deny IP fragmented packets, especially second and subsequent packets.

It extends the existing ACL command syntax with the *fragments* keyword for all Layer 3 rules applicable to all Layer protocols (permit/deny ip/tcp/udp/icmp).

- Both standard and extended ACLs support IP fragments.
- Second and subsequent fragments are allowed because a Layer 4 rule cannot be applied to these fragments. If the packet is to be denied eventually, the first fragment would be denied and hence the packet as a whole cannot be reassembled.
- Implementing the required rules uses a significant number of CAM entries per TCP/UDP entry.
- For IP ACL, the system always applies implicit deny. You do not have to configure it.
- For IP ACL, the system applies implicit permit for second and subsequent fragment prior to the implicit deny.
- If you configure an *explicit* deny, the second and subsequent fragments do not hit the implicit permit rule for fragments.

## IP Fragments ACL Examples

The following examples show how you can use ACL commands with the `fragment` keyword to filter fragmented packets.

### Example of Permitting All Packets on an Interface

### Example of Denying Second and Subsequent Fragments

The following configuration permits all packets (both fragmented and non-fragmented) with destination IP 10.1.1.1. The second rule does not get hit at all.

```
Dell(conf)#ip access-list extended ABC
Dell(conf-ext-nacl)#permit ip any 10.1.1.1/32
Dell(conf-ext-nacl)#deny ip any 10.1.1.1./32 fragments
Dell(conf-ext-nacl)
```

To deny the second/subsequent fragments, use the same rules in a different order. These ACLs deny all second and subsequent fragments with destination IP 10.1.1.1 but permit the first fragment and non-fragmented packets with destination IP 10.1.1.1.

```
Dell(conf)#ip access-list extended ABC
Dell(conf-ext-nacl)#deny ip any 10.1.1.1/32 fragments
Dell(conf-ext-nacl)#permit ip any 10.1.1.1/32
Dell(conf-ext-nacl)
```

## Layer 4 ACL Rules Examples

The following examples show the ACL commands for Layer 4 packet filtering.

When configuring ACLs with the `fragments` keyword, be aware of the following.

When an ACL filters packets, it looks at the fragment offset (FO) to determine whether it is a fragment.

- FO = 0 means it is either the first fragment or the packet is a non-fragment.
- FO > 0 means it is dealing with the fragments of the original packet.

### Permit an ACL line with L3 information only, and the `fragments` keyword is present:

If a packet's L3 information matches the L3 information in the ACL line, the packet's FO is checked.

- If a packet's FO > 0, the packet is permitted.
- If a packet's FO = 0, the next ACL entry is processed.

### Deny ACL line with L3 information only, and the `fragments` keyword is present:

If a packet's L3 information does match the L3 information in the ACL line, the packet's FO is checked.

- If a packet's FO > 0, the packet is denied.
- If a packet's FO = 0, the next ACL line is processed.

## Example of Layer 4 ACL Rules

### Example of TCP Packets

In this first example, fragments or non-fragmented TCP packets from 10.1.1.1 with TCP destination port equal to 24 are permitted. All other fragments are denied.

```
Dell(conf)#ip access-list extended ABC
Dell(conf-ext-nacl)#permit tcp host 10.1.1.1 any eq 24
Dell(conf-ext-nacl)#deny ip any any fragment
Dell(conf-ext-nacl)
```

In the following example, TCP packets that are first fragments or non-fragmented from host 10.1.1.1 with TCP destination port equal to 24 are permitted. Additionally, all TCP non-first fragments from host 10.1.1.1 are permitted. All other IP packets that are non-first fragments are denied.

```
Dell(conf)#ip access-list extended ABC
Dell(conf-ext-nacl)#permit tcp host 10.1.1.1 any eq 24
Dell(conf-ext-nacl)#permit tcp host 10.1.1.1 any fragment
Dell(conf-ext-nacl)#deny ip any any fragment
Dell(conf-ext-nacl)
```

## Configure a Standard IP ACL

To configure an ACL, use commands in IP ACCESS LIST mode and INTERFACE mode.

For a complete list of all the commands related to IP ACLs, refer to the *Dell Networking OS Command Line Interface Reference Guide*. To set up extended ACLs, refer to [Configure an Extended IP ACL](#).

A standard IP ACL uses the source IP address as its match criterion.

1. Enter IP ACCESS LIST mode by naming a standard IP access list.


CONFIGURATION mode

```
ip access-list standard access-listname
```

2. Configure a drop or forward filter.

CONFIG-STD-NACL mode

```
seq sequence-number {deny | permit} {source [mask] | any | host ip-address}
[count [byte]] [order] [fragments]
```

 **NOTE:** When assigning sequence numbers to filters, keep in mind that you might need to insert a new filter. To prevent reconfiguring multiple filters, assign sequence numbers in multiples of five.

When you use the `log` keyword, the CP logs details about the packets that match. Depending on how many packets match the log entry and at what rate, the CP may become busy as it has to log these packets' details.

To view the rules of a particular ACL configured on a particular interface, use the `show ip accounting access-list ACL-name interface interface` command in EXEC Privilege mode.

### Example of Viewing the Rules of a Specific ACL on an Interface

#### Example of the `seq` Command to Order Filters

```
Dell#show ip accounting access-list ToOspf interface gig 1/6
Standard IP access list ToOspf
```

```

seq 5 deny any
seq 10 deny 10.2.0.0 /16
seq 15 deny 10.3.0.0 /16
seq 20 deny 10.4.0.0 /16
seq 25 deny 10.5.0.0 /16
seq 30 deny 10.6.0.0 /16
seq 35 deny 10.7.0.0 /16
seq 40 deny 10.8.0.0 /16
seq 45 deny 10.9.0.0 /16
seq 50 deny 10.10.0.0 /16
Dell#

```

The following example shows how the `seq` command orders the filters according to the sequence number assigned. In the example, filter 25 was configured before filter 15, but the `show config` command displays the filters in the correct order.

```

Dell(conf-std-nacl)#seq 25 deny ip host 10.5.0.0 any
Dell(conf-std-nacl)#seq 15 permit tcp 10.3.0.0 /16 any
Dell(conf-std-nacl)#show config
!
ip access-list standard dilling
 seq 15 permit tcp 10.3.0.0/16 any
Dell(conf-std-nacl)#

```

To delete a filter, use the `no seq sequence-number` command in IP ACCESS LIST mode.

## Configuring a Standard IP ACL Filter

If you are creating a standard ACL with only one or two filters, you can let the system assign a sequence number based on the order in which the filters are configured. The software assigns filters in multiples of five.

1. Configure a standard IP ACL and assign it a unique name.

CONFIGURATION mode

```
ip access-list standard access-list-name
```

2. Configure a drop or forward IP ACL filter.

CONFIG-STD-NACL mode

```
{deny | permit} {source [mask] | any | host ip-address} [count [byte]]
[order] [fragments]
```

The following example shows a standard IP ACL in which the system assigns the sequence numbers. The filters were assigned sequence numbers based on the order in which they were configured (for example, the first filter was given the lowest sequence number). The `show config` command in IP ACCESS LIST mode displays the two filters with the sequence numbers 5 and 10.

### Example of Viewing Filter Sequence for a Specified Standard ACL

### Example of Viewing Standard ACL Filter Sequence for an Interface

```

Dell(config-route-map)#ip access standard kigali
Dell(config-std-nacl)#permit 10.1.0.0/16
Dell(config-std-nacl)#show config
!

```

```
ip access-list standard kigali
 seq 5 permit 10.1.0.0/16
Dell(config-std-nacl)#
```

To view all configured IP ACLs, use the `show ip accounting access-list` command in EXEC Privilege mode.

```
Dell#show ip accounting access example interface gig 4/12
Extended IP access list example
 seq 15 deny udp any any eq 111
 seq 20 deny udp any any eq 2049
 seq 25 deny udp any any eq 31337
 seq 30 deny tcp any any range 12345 12346
 seq 35 permit udp host 10.21.126.225 10.4.5.0 /28
 seq 40 permit udp host 10.21.126.226 10.4.5.0 /28
 seq 45 permit udp 10.8.0.0 /16 10.50.188.118 /31 range 1812 1813
 seq 50 permit tcp 10.8.0.0 /16 10.50.188.118 /31 eq 49
 seq 55 permit udp 10.15.1.0 /24 10.50.188.118 /31 range 1812 1813
```

To delete a filter, enter the `show config` command in IP ACCESS LIST mode and locate the sequence number of the filter you want to delete. Then use the `no seq sequence-number` command in IP ACCESS LIST mode.

## Configure an Extended IP ACL

Extended IP ACLs filter on source and destination IP addresses, IP host addresses, TCP addresses, TCP host addresses, UDP addresses, and UDP host addresses.

Because traffic passes through the filter in the order of the filter's sequence, you can configure the extended IP ACL by first entering IP ACCESS LIST mode and then assigning a sequence number to the filter.

## Configuring Filters with a Sequence Number

To configure filters with a sequence number, use the following commands.

1. Enter IP ACCESS LIST mode by creating an extended IP ACL.

CONFIGURATION mode

```
ip access-list extended access-list-name
```


2. Configure a drop or forward filter.

CONFIG-EXT-NAACL mode

```
seq sequence-number {deny | permit} {ip-protocol-number | icmp | ip | tcp |
udp} {source mask | any | host ip-address} {destination mask | any | host
ip-address} [operator port [port]] [count [byte]] [order] [fragments]
```

When you create the filters with a specific sequence number, you can create the filters in any order and the filters are placed in the correct order.



 **NOTE:** When assigning sequence numbers to filters, you might need to insert a new filter. To prevent reconfiguring multiple filters, assign sequence numbers in multiples of five or another number.

The following examples shows how the `seq` command orders the filters according to the sequence number assigned. In the example, filter 15 was configured before filter 5, but the `show config` command displays the filters in the correct order.

```
Dell(conf-ext-nacl)#seq 15 deny ip host 112.45.0.0 any
Dell(conf-ext-nacl)#seq 5 permit tcp 12.1.3.45 255.255.0.0 any
Dell(conf-ext-nacl)#show config
!
ip access-list extended dilling
 seq 5 permit tcp 12.1.0.0 255.255.0.0 any
 seq 15 deny ip host 112.45.0.0 any
Dell(conf-ext-nacl)#
```

## Configuring Filters Without a Sequence Number

If you are creating an extended ACL with only one or two filters, you can let the system assign a sequence number based on the order in which the filters are configured. The system assigns filters in multiples of five.

To configure a filter for an extended IP ACL without a specified sequence number, use any or all of the following commands:

- Configure a deny or permit filter to examine IP packets.

CONFIG-EXT-NACL mode

```
{deny | permit} {source mask | any | host ip-address} [count [byte]] [order]
[fragments]
```

- Configure a deny or permit filter to examine TCP packets.

CONFIG-EXT-NACL mode

```
{deny | permit} tcp {source mask} | any | host ip-address} [count [byte]]
[order] [fragments]
```

- Configure a deny or permit filter to examine UDP packets.

CONFIG-EXT-NACL mode

```
{deny | permit} udp {source mask | any | host ip-address} [count [byte]]
[order] [fragments]
```

The following example shows an extended IP ACL in which the sequence numbers were assigned by the software. The filters were assigned sequence numbers based on the order in which they were configured (for example, the first filter was given the lowest sequence number). The `show config` command in IP ACCESS LIST mode displays the two filters with the sequence numbers 5 and 10.

### Example of Viewing Filter Sequence for a Specified Extended ACL

```
Dell(config-ext-nacl)#deny tcp host 123.55.34.0 any
Dell(config-ext-nacl)#permit udp 154.44.123.34 0.0.255.255 host 34.6.0.0
Dell(config-ext-nacl)#show config
!
ip access-list extended nimule
 seq 5 deny tcp host 123.55.34.0 any
```

```
seq 10 permit udp 154.44.0.0 0.0.255.255 host 34.6.0.0
Dell(config-ext-nacl)#
```

To view all configured IP ACLs and the number of packets processed through the ACL, use the `show ip accounting access-list` command in EXEC Privilege mode, as shown in the first example in [Configuring a Standard IP ACL Filter](#).

## Established Flag

To obtain the functionality of **est**, use the following ACLs:

- `permit tcp any any rst`
- `permit tcp any any ack`

## Configure Layer 2 and Layer 3 ACLs

Both Layer 2 and Layer 3 ACLs may be configured on an interface in Layer 2 mode.

If both L2 and L3 ACLs are applied to an interface, the following rules apply:

- When the system routes the packets, only the L3 ACL governs them because they are not filtered against an L2 ACL.
- When the system switches the packets, first the L3 ACL filters them, then the L2 ACL filters them.
- When the system switches the packets, the egress L3 ACL does not filter the packet.


For the following features, if you enable counters on rules that have already been configured and a new rule is either inserted or prepended, all the existing counters are reset:

- L2 ingress access list
- L3 egress access list
- L2 egress access list
- L3 ingress access list

If a rule is simply appended, existing counters are not affected.

**Table 4. L2 and L3 Filtering on Switched Packets**

| L2 ACL Behavior | L3 ACL Behavior | Decision on Targeted Traffic |
|-----------------|-----------------|------------------------------|
| Deny            | Deny            | L3 ACL denies.               |
| Deny            | Permit          | L3 ACL permits.              |
| Permit          | Deny            | L3 ACL denies.               |
| Permit          | Permit          | L3 ACL permits.              |

 **NOTE:** If you configure an interface as a vlan-stack access port, only the L2 ACL filters the packets. The L3 ACL applied to such a port does not affect traffic. That is, existing rules for other features (such as trace-list, policy-based routing [PBR], and QoS) are applied to the permitted traffic.

For information about MAC ACLs, refer to [Layer 2](#).

## Assign an IP ACL to an Interface

To pass traffic through a configured IP ACL, assign that ACL to a physical interface, a port channel interface, or a VLAN.

The IP ACL is applied to all traffic entering a physical or port channel interface and the traffic is either forwarded or dropped depending on the criteria and actions specified in the ACL.

The same ACL may be applied to different interfaces and that changes its functionality. For example, you can take ACL "ABCD" and apply it using the `in` keyword and it becomes an ingress access list. If you apply the same ACL using the `out` keyword, it becomes an egress access list.

For more information about Layer-3 interfaces, refer to [Interfaces](#).

## Applying an IP ACL

To apply an IP ACL (standard or extended) to a physical or port channel interface, use the following commands.

1. Enter the interface number.

CONFIGURATION mode

```
interface interface slot/port
```

2. Configure an IP address for the interface, placing it in Layer-3 mode.

INTERFACE mode

```
ip address ip-address
```

3. Apply an IP ACL to traffic entering or exiting an interface.

INTERFACE mode

```
ip access-group access-list-name {in | out} [implicit-permit] [vlan vlan-range]
```



**NOTE:** The number of entries allowed per ACL is hardware-dependent. For detailed specification about entries allowed per ACL, refer to your line card documentation.

4. Apply rules to the new ACL.

INTERFACE mode

```
ip access-list [standard | extended] name
```

To view which IP ACL is applied to an interface, use the `show config` command in INTERFACE mode, or use the `show running-config` command in EXEC mode.

### Example of Viewing ACLs Applied to an Interface

```
Dell(conf-if)#show conf
!
interface GigabitEthernet 0/0
 ip address 10.2.1.100 255.255.255.0
 ip access-group nimule in
```

```
no shutdown
Dell(conf-if)#
```

To filter traffic on Telnet sessions, use only standard ACLs in the `access-class` command.

## Counting ACL Hits

You can view the number of packets matching the ACL by using the `count` option when creating ACL entries.

In the MXL switch, you can configure either `count (packets)` or `count (bytes)`. However, for an ACL with multiple rules, you can configure some ACLs with `count (packets)` and others as `count (bytes)` at any given time.

1. Create an ACL that uses rules with the `count` option. Refer to [Configuring a Standard IP ACL Filter](#).
2. Apply the ACL as an inbound or outbound ACL on an interface. Refer to [Assign an IP ACL to an Interface](#).
3. View the number of packets matching the ACL.

```
EXEC Privilege mode
```

```
show ip accounting access-list
```

## Configure Ingress ACLs

Ingress ACLs are applied to interfaces and to traffic entering the system.

These system-wide ACLs eliminate the need to apply ACLs onto each interface and achieves the same results. By localizing target traffic, it is a simpler implementation.

To create an ingress ACL, use the `ip access-group` command in EXEC Privilege mode. The example shows applying the ACL, rules to the newly created access group, and viewing the access list.

### Example of Applying ACL Rules to Ingress Traffic and Viewing ACL Configuration

To specify ingress, use the `in` keyword. Begin applying rules to the ACL with the `ip access-list extended abcd` command. To view the access-list, use the `show` command.

```
Dell(conf)#interface tengig 0/0
Dell(conf-if-tengig0/0)#ip access-group abcd in
Dell(conf-if-tengig0/0)#show config
!
tengigetherenet 0/0
no ip address
ip access-group abcd in
no shutdown
Dell(conf-if-tengig0/0)#end
Dell#configure terminal
Dell(conf)#ip access-list extended abcd
Dell(conf-ext-nacl)#permit tcp any any
Dell(conf-ext-nacl)#deny icmp any any
Dell(conf-ext-nacl)#permit 1.1.1.2
Dell(conf-ext-nacl)#end
Dell#show ip accounting access-list
!
Extended Ingress IP access list abcd on tengigetherenet 0/0
```

```
seq 5 permit tcp any any
seq 10 deny icmp any any
seq 15 permit 1.1.1.2
```

## Configure Egress ACLs

Configuring egress ACLs onto physical interfaces protects the system infrastructure from attack — malicious and incidental — by explicitly allowing only authorized traffic.

These system-wide ACLs eliminate the need to apply ACLs onto each interface and achieves the same results. By localizing target traffic, it is a simpler implementation.

To restrict egress traffic, use an egress ACL. For example, when a direct operating system (DOS) attack traffic is isolated to a specific interface, you can apply an egress ACL to block the flow from the exiting the box, thus protecting downstream devices.

To create an egress ACL, use the `ip access-group` command in EXEC Privilege mode. The example shows viewing the configuration, applying rules to the newly created access group, and viewing the access list.

### Example of Applying ACL Rules to Egress Traffic and Viewing ACL Configuration

To specify ingress, use the `out` keyword. Begin applying rules to the ACL with the `ip access-list extended abcd` command. To view the access-list, use the `show` command.

```
Dell(conf)#interface tengig 0/0
Dell(conf-if-tengig0/0)#ip access-group abcd out
Dell(conf-if-tengig0/0)#show config
!
tengigetherenet 0/0
 no ip address
 ip access-group abcd out
 no shutdown
Dell(conf-if-tengig0/0)#end
Dell#configure terminal
Dell(conf)#ip access-list extended abcd
Dell(conf-ext-nacl)#permit tcp any any
Dell(conf-ext-nacl)#deny icmp any any
Dell(conf-ext-nacl)#permit 1.1.1.2
Dell(conf-ext-nacl)#end
Dell#show ip accounting access-list
!
Extended Ingress IP access list abcd on tengigetherenet 0/0
 seq 5 permit tcp any any
 seq 10 deny icmp any any
 seq 15 permit 1.1.1.2
```

## Applying Egress Layer 3 ACLs (Control-Plane)

By default, packets originated from the system are not filtered by egress ACLs.

For example, if you initiate a ping session from the system and apply an egress ACL to block this type of traffic on the interface, the ACL does not affect that ping traffic. The Control Plane Egress Layer 3 ACL feature enhances IP reachability debugging by implementing control-plane ACLs for CPU-generated and

CPU-forwarded traffic. Using permit rules with the `count` option, you can track on a per-flow basis whether CPU-generated and CPU-forwarded packets were transmitted successfully.

1. Apply Egress ACLs to IPv4 system traffic.

CONFIGURATION mode

```
ip control-plane [egress filter]
```

2. Create a Layer 3 ACL using permit rules with the `count` option to describe the desired CPU traffic.

CONFIG-NACL mode

```
permit ip {source mask | any | host ip-address} {destination mask | any |
host ip-address} count
```

**Dell Networking OS Behavior:** Virtual router redundancy protocol (VRRP) hellos and internet group management protocol (IGMP) packets are not affected when you enable egress ACL filtering for CPU traffic. Packets sent by the CPU with the source address as the VRRP virtual IP address have the interface MAC address instead of VRRP virtual MAC address.

## IP Prefix Lists

IP prefix lists control routing policy.

An IP prefix list is a series of sequential filters that contain a matching criterion (examine IP route prefix) and an action (permit or deny) to process routes. The filters are processed in sequence so that if a route prefix does not match the criterion in the first filter, the second filter (if configured) is applied. When the route prefix matches a filter, the system drops or forwards the packet based on the filter's designated action. If the route prefix does not match any of the filters in the prefix list, the route is dropped (that is, implicit deny).

A route prefix is an IP address pattern that matches on bits within the IP address. The format of a route prefix is A.B.C.D/X where A.B.C.D is a dotted-decimal address and /X is the number of bits that should be matched of the dotted decimal address. For example, in 112.24.0.0/16, the first 16 bits of the address 112.24.0.0 match all addresses between 112.24.0.0 to 112.24.255.255.

The following examples show permit or deny filters for specific routes using the `le` and `ge` parameters, where x.x.x.x/x represents a route prefix:


- To deny only /8 prefixes, enter `deny x.x.x.x/x ge 8 le 8`.
- To permit routes with the mask greater than /8 but less than /12, enter `permit x.x.x.x/x ge 8`.
- To deny routes with a mask less than /24, enter `deny x.x.x.x/x le 24`.
- To permit routes with a mask greater than /20, enter `permit x.x.x.x/x ge 20`.

The following rules apply to prefix lists:

- A prefix list without any permit or deny filters allows all routes.
- An "implicit deny" is assumed (that is, the route is dropped) for all route prefixes that do not match a permit or deny filter in a configured prefix list.
- After a route matches a filter, the filter's action is applied. No additional filters are applied to the route.

## Implementation Information

In the Dell Networking OS, prefix lists are used in processing routes for routing protocols (for example, router information protocol [RIP], open shortest path first [OSPF], and border gateway protocol [BGP]).

 **NOTE:** The MXL Switch platform does not support all protocols. It is important to know which protocol you are supporting prior to implementing prefix lists.

## Configuration Task List for Prefix Lists

To configure a prefix list, use commands in PREFIX LIST, ROUTER RIP, ROUTER OSPF, and ROUTER BGP modes.

Create the prefix list in PREFIX LIST mode and assign that list to commands in ROUTER RIP, ROUTER OSPF and ROUTER BGP modes.

The following list includes the configuration tasks for prefix lists, as described in the following sections.

- Configuring a prefix list
- Use a prefix list for route redistribution

For a complete listing of all commands related to prefix lists, refer to the *Dell Networking OS Command Line Interface Reference Guide*.

## Creating a Prefix List

To create a prefix list, use the following commands.

1. Create a prefix list and assign it a unique name.  
You are in PREFIX LIST mode.

```
CONFIGURATION mode
```

```
ip prefix-list prefix-name
```

2. Create a prefix list with a sequence number and a deny or permit action.  
CONFIG-NPREFIXL mode

```
seq sequence-number {deny | permit} ip-prefix [ge min-prefix-length] [le max-prefix-length]
```

The optional parameters are:


- *ge min-prefix-length*: the minimum prefix length to match (from 0 to 32).
- *le max-prefix-length*: the maximum prefix length to match (from 0 to 32).

### Example of Assigning Sequence Numbers to Filters

If you want to forward all routes that do not match the prefix list criteria, configure a prefix list filter to permit all routes (`permit 0.0.0.0/0 le 32`). The “permit all” filter must be the last filter in your prefix list. To permit the default route only, enter `permit 0.0.0.0/0`.

The following example shows how the `seq` command orders the filters according to the sequence number assigned. In the example, filter 20 was configured before filter 15 and 12, but the `show config` command displays the filters in the correct order.

```
Dell(conf-nprefixl)#seq 20 permit 0.0.0.0/0 le 32
Dell(conf-nprefixl)#seq 12 deny 134.23.0.0 /16
Dell(conf-nprefixl)#seq 15 deny 120.23.14.0 /8 le 16
Dell(conf-nprefixl)#show config
!
ip prefix-list juba
 seq 12 deny 134.23.0.0/16
 seq 15 deny 120.0.0.0/8 le 16
 seq 20 permit 0.0.0.0/0 le 32
Dell(conf-nprefixl)#
```

 **NOTE:** The last line in the prefix list Juba contains a “permit all” statement. By including this line in a prefix list, you specify that all routes not matching any criteria in the prefix list are forwarded.

To delete a filter, use the `no seq sequence-number` command in PREFIX LIST mode.

If you are creating a standard prefix list with only one or two filters, you can let the system assign a sequence number based on the order in which the filters are configured. The system assigns filters in multiples of five.

## Creating a Prefix List Without a Sequence Number

To create a filter without a specified sequence number, use the following commands.

1. Create a prefix list and assign it a unique name.

CONFIGURATION mode

```
ip prefix-list prefix-name
```

2. Create a prefix list filter with a deny or permit action.

CONFIG-NPREFIXL mode

```
{deny | permit} ip-prefix [ge min-prefix-length] [le max-prefix-length]
```

The optional parameters are:

- `ge min-prefix-length`: is the minimum prefix length to be matched (from 0 to 32).
- `le max-prefix-length`: is the maximum prefix length to be matched (from 0 to 32).

### Example of Creating a Filter with a Dell Networking OS-Assigned Sequence Numbers

The example shows a prefix list in which the sequence numbers were assigned by the software. The filters were assigned sequence numbers based on the order in which they were configured (for example, the first filter was given the lowest sequence number). The `show config` command in PREFIX LIST mode displays the two filters with the sequence numbers 5 and 10.

```
Dell(conf-nprefixl)#permit 123.23.0.0 /16
Dell(conf-nprefixl)#deny 133.24.56.0 /8
Dell(conf-nprefixl)#show conf
!
ip prefix-list awe
 seq 5 permit 123.23.0.0/16
```



```
seq 10 deny 133.0.0.0/8
Dell(conf-nprefix1)#
```

To delete a filter, enter the `show config` command in PREFIX LIST mode and locate the sequence number of the filter you want to delete, then use the `no seq sequence-number` command in PREFIX LIST mode.

## Viewing Prefix Lists

To view all configured prefix lists, use the following commands.

- Show detailed information about configured prefix lists.

EXEC Privilege mode

```
show ip prefix-list detail [prefix-name]
```

- Show a table of summarized information about configured Prefix lists.

EXEC Privilege mode

```
show ip prefix-list summary [prefix-name]
```

### Example of the `show ip prefix-list detail` Command

### Example of the `show ip prefix-list summary` Command

```
Dell>show ip prefix detail
Prefix-list with the last deletion/insertion: filter_ospf
ip prefix-list filter_in:
count: 3, range entries: 3, sequences: 5 - 10
 seq 5 deny 1.102.0.0/16 le 32 (hit count: 0)
 seq 6 deny 2.1.0.0/16 ge 23 (hit count: 0)
 seq 10 permit 0.0.0.0/0 le 32 (hit count: 0)
ip prefix-list filter_ospf:
count: 4, range entries: 1, sequences: 5 - 10
 seq 5 deny 100.100.1.0/24 (hit count: 0)
 seq 6 deny 200.200.1.0/24 (hit count: 0)
 seq 7 deny 200.200.2.0/24 (hit count: 0)
 seq 10 permit 0.0.0.0/0 le 32 (hit count: 0)
```

```
Dell>
Dell>show ip prefix summary
Prefix-list with the last deletion/insertion: filter_ospf
ip prefix-list filter_in:
count: 3, range entries: 3, sequences: 5 - 10
ip prefix-list filter_ospf:
count: 4, range entries: 1, sequences: 5 - 10
Dell>
```

## Applying a Prefix List for Route Redistribution

To pass traffic through a configured prefix list, use the prefix list in a `route redistribution` command.

Apply the prefix list to all traffic redistributed into the routing process. The traffic is either forwarded or dropped, depending on the criteria and actions specified in the prefix list.

To apply a filter to routes in RIP, use the following commands.

- Enter RIP mode.

CONFIGURATION mode

```
router rip
```

- Apply a configured prefix list to incoming routes. You can specify an interface. If you enter the name of a nonexistent prefix list, all routes are forwarded.

CONFIG-ROUTER-RIP mode

```
distribute-list prefix-list-name in [interface]
```

- Apply a configured prefix list to outgoing routes. You can specify an interface or type of route. If you enter the name of a non-existent prefix list, all routes are forwarded.

CONFIG-ROUTER-RIP mode

```
distribute-list prefix-list-name out [interface | connected | static | ospf]
```

### Example of Viewing Configured Prefix Lists (ROUTER RIP mode)

To view the configuration, use the `show config` command in ROUTER RIP mode, or the `show running-config rip` command in EXEC mode.

```
Dell(conf-router_rip)#show config
!
router rip
 distribute-list prefix juba out
 network 10.0.0.0
Dell(conf-router_rip)#router ospf 34
```

## Applying a Filter to a Prefix List (OSPF)

To apply a filter to routes in open shortest path first (OSPF), use the following commands.

- Enter OSPF mode.

CONFIGURATION mode

```
router ospf
```

- Apply a configured prefix list to incoming routes. You can specify an interface. If you enter the name of a non-existent prefix list, all routes are forwarded.

CONFIG-ROUTER-OSPF mode

```
distribute-list prefix-list-name in [interface]
```

- Apply a configured prefix list to incoming routes. You can specify which type of routes are affected. If you enter the name of a non-existent prefix list, all routes are forwarded.

CONFIG-ROUTER-OSPF mode

```
distribute-list prefix-list-name out [connected | rip | static]
```

### Example of Viewing Configured Prefix Lists (ROUTER OSPF mode)

To view the configuration, use the `show config` command in ROUTER OSPF mode, or the `show running-config ospf` command in EXEC mode.

```
Dell(conf-router_ospf)#show config
!
router ospf 34
 network 10.2.1.1 255.255.255.255 area 0.0.0.1
 distribute-list prefix awe in
Dell(conf-router_ospf)#
```


## ACL Resequencing

ACL resequencing allows you to re-number the rules and remarks in an access or prefix list.

The placement of rules within the list is critical because packets are matched against rules in sequential order. To order new rules using the current numbering scheme, use resequencing whenever there is no opportunity.

For example, the following table contains some rules that are numbered in increments of 1. You cannot place new rules between these packets, so apply resequencing to create numbering space, as shown in the second table. In the same example, apply resequencing if more than two rules must be placed between rules 7 and 10.

You can resequence IPv4 ACLs, prefixes, and MAC ACLs. No CAM writes happen as a result of resequencing, so there is no packet loss; the behavior is similar Hot-lock ACLs.

 **NOTE:** ACL resequencing does not affect the rules, remarks, or order in which they are applied. Resequencing merely rennumbers the rules so that you can place new rules within the list as needed.

**Table 5. ACL Resequencing**

| Rules                      | Resequencing                   |
|----------------------------|--------------------------------|
| Rules Before Resequencing: | seq 5 permit any host 1.1.1.1  |
|                            | seq 6 permit any host 1.1.1.2  |
|                            | seq 7 permit any host 1.1.1.3  |
|                            | seq 10 permit any host 1.1.1.4 |
| Rules After Resequencing:  | seq 5 permit any host 1.1.1.1  |
|                            | seq 10 permit any host 1.1.1.2 |
|                            | seq 15 permit any host 1.1.1.3 |
|                            | seq 20 permit any host 1.1.1.4 |

## Resequencing an ACL or Prefix List

Resequencing is available for IPv4 ACLs, prefix lists, and MAC ACLs.

To resequence an ACL or prefix list, use the following commands. You must specify the list name, starting number, and increment when using these commands.

- Resequence an IPv4 or MAC ACL.

EXEC mode

```
resequence access-list {ipv4 | mac} {access-list-name StartingSeqNum Step-to-Increment}
```

- Resequence an IPv4 prefix-list.

EXEC mode

```
resequence prefix-list {ipv4} {prefix-list-name StartingSeqNum Step-to-Increment}
```

### Example of Resequencing ACLs When Remarks and Rules Have the Same Number

### Example of Resequencing ACLs When Remarks and Rules Have Different Numbers

The example shows the resequencing of an IPv4 access-list beginning with the number 2 and incrementing by 2.

Remarks and rules that originally have the same sequence number have the same sequence number after you apply the `resequence` command.

```
Dell(config-ext-nacl)# show config
!
ip access-list extended test
remark 4 XYZ
remark 5 this remark corresponds to permit any host 1.1.1.1
seq 5 permit ip any host 1.1.1.1
remark 9 ABC
remark 10 this remark corresponds to permit ip any host 1.1.1.2
seq 10 permit ip any host 1.1.1.2
seq 15 permit ip any host 1.1.1.3
seq 20 permit ip any host 1.1.1.4
Dell# end
Dell# resequence access-list ipv4 test 2 2
Dell# show running-config acl
!
ip access-list extended test
remark 2 XYZ
remark 4 this remark corresponds to permit any host 1.1.1.1
seq 4 permit ip any host 1.1.1.1
remark 6 this remark has no corresponding rule
remark 8 this remark corresponds to permit ip any host 1.1.1.2
seq 8 permit ip any host 1.1.1.2
seq 10 permit ip any host 1.1.1.3
seq 12 permit ip any host 1.1.1.4
```

Remarks that do not have a corresponding rule are incremented as a rule. These two mechanisms allow remarks to retain their original position in the list. The following example shows remark 10 corresponding to rule 10 and as such, they have the same number before and after the command is entered. Remark 4 is incremented as a rule, and all rules have retained their original positions.

```
Dell(config-ext-nacl)# show config
!
ip access-list extended test
remark 4 XYZ
remark 5 this remark corresponds to permit any host 1.1.1.1
seq 5 permit ip any host 1.1.1.1
remark 9 ABC
remark 10 this remark corresponds to permit ip any host 1.1.1.2
```

```

seq 10 permit ip any host 1.1.1.2
seq 15 permit ip any host 1.1.1.3
seq 20 permit ip any host 1.1.1.4
Dell# end
Dell# resequence access-list ipv4 test 2 2
Dell# show running-config acl
!
ip access-list extended test
remark 2 XYZ
remark 4 this remark corresponds to permit any host 1.1.1.1
seq 4 permit ip any host 1.1.1.1
remark 6 this remark has no corresponding rule
remark 8 this remark corresponds to permit ip any host 1.1.1.2
seq 8 permit ip any host 1.1.1.2
seq 10 permit ip any host 1.1.1.3
seq 12 permit ip any host 1.1.1.4

```

## Route Maps

Similar to ACLs and prefix lists, route maps are composed of a series of commands that contain a matching criterion and an action; however, route maps can change the packets meeting the criterion.

ACLs and prefix lists can only drop or forward the packet or traffic. Route maps process routes for route redistribution. For example, a route map can be called to filter only specific routes and to add a metric.

Route maps also have an "implicit deny." Unlike ACLs and prefix lists; however, where the packet or traffic is dropped, in route maps, if a route does not match any of the route map conditions, the route is not redistributed.

### Implementation Information

The Dell Networking OS implementation of route maps allows route maps with the `no match` or `no set` commands. When there is `no match` command, all traffic matches the route map and the `set` command applies.

## Important Points to Remember

- For route-maps with more than one match clause:
  - Two or more match clauses within the same route-map sequence have the *same* match commands (though the values are different), matching a packet against these clauses is a logical OR operation.
  - Two or more match clauses within the same route-map sequence have *different* match commands, matching a packet against these clauses is a logical AND operation.
- If no match is found in a route-map sequence, the process moves to the next route-map sequence until a match is found, or there are no more sequences.
- When a match is found, the packet is forwarded and no more route-map sequences are processed.
  - If a continue clause is included in the route-map sequence, the next or a specified route-map sequence is processed after a match is found.

## Configuration Task List for Route Maps

Configure route maps in ROUTE-MAP mode and apply the maps in various commands in ROUTER RIP and ROUTER OSPF modes.

The following list includes the configuration tasks for route maps, as described in the following sections.

- Create a route map (mandatory)
- Configure route map filters (optional)
- Configure a route map for route redistribution (optional)
- Configure a route map for route tagging (optional)

## Creating a Route Map

Route maps, ACLs, and prefix lists are similar in composition because all three contain filters, but route map filters do not contain the permit and deny actions found in ACLs and prefix lists. Route map filters match certain routes and set or specify values.

To create a route map, use the following command.

- Create a route map and assign it a unique name. The optional `permit` and `deny` keywords are the action of the route map.

CONFIGURATION mode

```
route-map map-name [permit | deny] [sequence-number]
```

The default is **permit**.

The optional `seq` keyword allows you to assign a sequence number to the route map instance.

### Example of Viewing a Configured Route Map

### Example of Multiple Instances of a Route-Map

### Example of Deleting One Instance of a Route Map

### Example of Viewing All Instances of a Specified Route Map

The default action is **permit** and the default sequence number starts at **10**. When you use the keyword `deny` in configuring a route map, routes that meet the match filters are not redistributed.

To view the configuration, use the `show config` command in ROUTE-MAP mode.

```
Dell(config-route-map)#show config
!
route-map dilling permit 10
Dell(config-route-map)#
```

You can create multiple instances of this route map by using the `sequence number` option to place the route maps in the correct order. The system processes the route maps with the lowest sequence number first. When a configured route map is applied to a command, such as `redistribute`, traffic passes

through all instances of that route map until a match is found. The following is an example with two instances of a route map.

```
Dell#show route-map
route-map zakho, permit, sequence 10
 Match clauses:
 Set clauses:
route-map zakho, permit, sequence 20
 Match clauses:
 interface TenGigabitEthernet 0/1
 Set clauses:
 tag 35
 level stub-area
Dell#
```

To delete all instances of that route map, use the `no route-map map-name` command. To delete just one instance, add the sequence number to the command syntax.

```
Dell(conf)#no route-map zakho 10
Dell(conf)#end
Dell#show route-map
route-map zakho, permit, sequence 20
 Match clauses:
 interface TenGigabitEthernet 0/1
 Set clauses:
 tag 35
 level stub-area
Dell#
```

The following example shows a route map with multiple instances. The `show config` command displays only the configuration of the current route map instance. To view all instances of a specific route map, use the `show route-map` command.

```
Dell#show route-map dilling
route-map dilling, permit, sequence 10
 Match clauses:
 Set clauses:
route-map dilling, permit, sequence 15
 Match clauses:
 interface Loopback 23
 Set clauses:
 tag 3444
Dell#
```

To delete a route map, use the `no route-map map-name` command in CONFIGURATION mode.

## Configure Route Map Filters

Within ROUTE-MAP mode, there are `match` and `set` commands.

- `match` commands search for a certain criterion in the routes.
- `set` commands change the characteristics of routes, either adding something or specifying a level.

When there are multiple `match` commands with the same parameter under one instance of `route-map`, the system does a match between all of those `match` commands. If there are multiple `match` commands

with different parameters, the system does a match ONLY if there is a match among ALL the `match` commands.

In the following example, there is a match if a route has any of the tag values specified in the `match` commands.

### Example of the `match` Command to Match Any of Several Values

#### Example of the `match` Command to Match All Specified Values

```
Dell(conf)#route-map force permit 10
Dell(config-route-map)#match tag 1000
Dell(config-route-map)#match tag 2000
Dell(config-route-map)#match tag 3000
```

In the next example, there is a match *only* if a route has *both* of the specified characteristics. In this example, there a match only if the route has a tag value of 1000 *and* a metric value of 2000.

Also, if there are different instances of the same route-map, then it's sufficient if a permit match happens in any instance of that route-map.

```
Dell(conf)#route-map force permit 10
Dell(config-route-map)#match tag 1000
Dell(config-route-map)#match metric 2000
```

In the following example, instance 10 permits the route having a tag value of 1000 and instances 20 and 30 deny the route having a tag value of 1000. In this scenario, the system scans all the instances of the route-map for any permit statement. If there is a match anywhere, the route is permitted. However, other instances of the route-map deny it.

### Example of the `match` Command to Permit and Deny Routes

```
Dell(conf)#route-map force permit 10
Dell(config-route-map)#match tag 1000

Dell(conf)#route-map force deny 20
Dell(config-route-map)#match tag 1000

Dell(conf)#route-map force deny 30
Dell(config-route-map)#match tag 1000
```

## Configuring Match Routes

To configure match criterion for a route map, use the following commands.

- Match routes whose next hop is a specific interface.

CONFIG-ROUTE-MAP mode

```
match interface interface
```



The parameters are:

- For a Loopback interface, enter the keyword `loopback` then a number between zero (0) and 16383.
- For a 10-Gigabit Ethernet interface, enter the keyword `tengigabitEthernet` then the slot/port information.
- For a VLAN, enter the keyword `vlan` then a number from 1 to 4094.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
- Match destination routes specified in a prefix list (IPv4).

CONFIG-ROUTE-MAP mode

```
match ip address prefix-list-name
```

- Match next-hop routes specified in a prefix list (IPv4).

CONFIG-ROUTE-MAP mode

```
match ip next-hop {access-list-name | prefix-list prefix-list-name}
```

- Match source routes specified in a prefix list (IPv4).

CONFIG-ROUTE-MAP mode

```
match ip route-source {access-list-name | prefix-list prefix-list-name}
```

- Match routes with a specific value.

CONFIG-ROUTE-MAP mode

```
match metric metric-value
```

- Match routes specified as internal or external to OSPF, ISIS level-1, ISIS level-2, or locally generated.

CONFIG-ROUTE-MAP mode

```
match route-type {external [type-1 | type-2] | internal | level-1 | level-2 | local }
```

- Match routes with a specific tag.

CONFIG-ROUTE-MAP mode

```
match tag tag-value
```

To create route map instances, use these commands. There is no limit to the number of `match` commands per route map, but the convention is to keep the number of match filters in a route map low. `Set` commands do not require a corresponding `match` command.

## Configuring Set Conditions

To configure a set condition, use the following commands.

- Generate a tag to be added to redistributed routes.

CONFIG-ROUTE-MAP mode

```
set automatic-tag
```

- Specify an OSPF area or ISIS level for redistributed routes.

CONFIG-ROUTE-MAP mode

```
set level {backbone | level-1 | level-1-2 | level-2 | stub-area}
```

- Specify a value for redistributed routes.

CONFIG-ROUTE-MAP mode

```
set metric {+ | - | metric-value}
```

- Specify an OSPF or ISIS type for redistributed routes.

CONFIG-ROUTE-MAP mode

```
set metric-type {external | internal | type-1 | type-2}
```

- Assign an IP address as the route's next hop.

CONFIG-ROUTE-MAP mode

```
set next-hop ip-address
```

- Specify a tag for the redistributed routes.

CONFIG-ROUTE-MAP mode

```
set tag tag-value
```

To create route map instances, use these commands. There is no limit to the number of `set` commands per route map, but the convention is to keep the number of set filters in a route map low. Set commands do not require a corresponding `match` command.


## Configure a Route Map for Route Redistribution

Route maps on their own cannot affect traffic and must be included in different commands to affect routing traffic. To apply a route map to traffic, you must call or include that route map in a command such as the `redistribute` or `default-information originate` commands in OSPF and BGP.

Route redistribution occurs when the system learns the advertising routes from static or directly connected routes or another routing protocol. Different protocols assign different values to redistributed routes to identify either the routes and their origins. The metric value is the most common attribute that is changed to properly redistribute other routes into a routing protocol. Other attributes that can be changed include the metric type (for example, external and internal route types in OSPF) and route tag. Use the `redistribute` command in OSPF, RIP, ISIS, and BGP to set some of these attributes for routes that are redistributed into those protocols.

Route maps add to that redistribution capability by allowing you to match specific routes and set or change more attributes when redistributing those routes.

In the following example, the `redistribute` command calls the route map `static ospf` to redistribute only certain static routes into OSPF. According to the route map `static ospf`, only routes that have a next hop of GigabitEthernet interface 0/0 and that have a metric of 255 are redistributed into the OSPF backbone area.

 **NOTE:** When re-distributing routes using route-maps, you must create the route-map defined in the `redistribute` command under the routing protocol. If you do not create a route-map, NO routes are redistributed.

### Example of Calling a Route Map to Redistribute Specified Routes

```
router ospf 34
 default-information originate metric-type 1
```

```

 redistribute static metric 20 metric-type 2 tag 0 route-map staticospf
!
route-map staticospf permit 10
 match interface GigabitEthernet 0/0
 match metric 255
 set level backbone

```

## Configure a Route Map for Route Tagging

One method for identifying routes from different routing protocols is to assign a tag to routes from that protocol.

As the route enters a different routing domain, it is tagged. The tag is passed along with the route as it passes through different routing protocols. You can use this tag when the route leaves a routing domain to redistribute those routes again.

In the following example, the `redistribute ospf` command with a route map is used in ROUTER RIP mode to apply a tag of 34 to all internal OSPF routes that are redistributed into RIP.

### Example of the `redistribute` Command Using a Route Tag

```


!
router rip
 redistribute ospf 34 metric 1 route-map torip
!
route-map torip permit 10
 match route-type internal
 set tag 34
!

```

## Continue Clause

Normally, when a match is found, set clauses are executed, and the packet is then forwarded; no more route-map modules are processed.

If you configure the `continue` command at the end of a module, the next module (or a specified module) is processed even after a match is found. The following example shows a continue clause at the end of a route-map module. In this example, if a match is found in the route-map “test” module 10, module 30 is processed.

 **NOTE:** If you configure the continue clause without specifying a module, the next sequential module is processed.

### Example of Using the `continue` Clause in a Route Map

```

!
route-map test permit 10
match commu comm-list1
set community 1:1 1:2 1:3
set as-path prepend 1 2 3 4 5
continue 30!

```

## Logging of ACL Processes

To assist in the administration and management of traffic that traverses the device after being validated by the configured ACLs, you can enable the generation of logs for access control list (ACL) processes. Although you can configure ACLs with the required permit or deny filters to provide access to the incoming packet or disallow access to a particular user, it is also necessary to monitor and examine the traffic that passes through the device. To evaluate network traffic that is subjected to ACLs, configure the logs to be triggered for ACL operations. This functionality is primarily needed for network supervision and maintenance activities of the handled subscriber traffic.

When ACL logging is configured, and a frame reaches an ACL-enabled interface and matches the ACL, a log is generated to indicate that the ACL entry matched the packet.

When you enable ACL log messages, at times, depending on the volume of traffic, it is possible that a large number of logs might be generated that can impact the system performance and efficiency. To avoid an overload of ACL logs from being recorded, you can configure the rate-limiting functionality. Specify the interval or frequency at which ACL logs must be triggered and also the threshold or limit for the maximum number of logs to be generated. If you do not specify the frequency at which ACL logs must be generated, a default interval of 5 minutes is used. Similarly, if you do not specify the threshold for ACL logs, a default threshold of 10 is used, where this value refers to the number of packets that are matched against an ACL .

A Layer 2 or Layer 3 ACL contains a set of defined rules that are saved as flow processor (FP) entries. When you enable ACL logging for a particular ACL rule, a set of specific ACL rules translate to a set of FP entries. You can enable logging separately for each of these FP entries, which relate to each of the ACL entries configured in an ACL. Dell Networking OS saves a table that maps each ACL entry that matches the ACL name on the received packet, sequence number of the rule, and the interface index in the database. When the configured maximum threshold has exceeded, log generation stops. When the interval at which ACL logs are configured to be recorded expires, a fresh interval timer starts and the packet count for that new interval commences from zero. If ACL logging was stopped previously because the configured threshold has exceeded, it is reenabled for this new interval.

The ACL application sends the ACL logging configuration information and other details, such as the action, sequence number, and the ACL parameters that pertain to that ACL entry. The ACL service collects the ACL log and records the following attributes per log message.

- For non-IP packets, the ACL name, sequence number, ACL action (permit or deny), source and destination MAC addresses, EtherType, and ingress interface are the logged attributes.
- For IP Packets, the ACL name, sequence number, ACL action (permit or deny), source and destination MAC addresses, source and destination IP addresses, and the transport layer protocol used are the logged attributes.
- For IP packets that contain the transport layer protocol as Transmission Control Protocol (TCP) or User Datagram Protocol (UDP), the ACL name, sequence number, ACL action (permit or deny), source and destination MAC addresses, source and destination IP addresses, and the source and destination ports (Layer 4 parameters) are also recorded.

If the packet contains an unidentified EtherType or transport layer protocol, the values for these parameters are saved as Unknown in the log message. If you also enable the logging of the count of

packets in the ACL entry, and if the logging is deactivated in a specific interval because the threshold has exceeded, the count of packets that exceeded the logging threshold value during that interval is recorded when the subsequent log record (in the next interval) is generated for that ACL entry.

## Guidelines for Configuring ACL Logging

Keep the following points in mind when you configure logging of ACL activities:

- During initialization, the ACL logging application tags the ACL rule indices for which a match condition exists as being in-use, which ensures that the same rule indices are not reused by ACL logging again.
- The ACL configuration information that the ACL logging application receives from the ACL manager causes the allocation and clearance of the match rule number. A unique match rule number is created for the combination of each ACL entry, sequence number, and interface parameters.
- A separate set of match indices is preserved by the ACL logging application for the permit and deny actions. Depending on the action of an ACL entry, the corresponding match index is allocated from the particular set that is maintained for permit and deny actions.
- A maximum of 125 ACL entries with permit action can be logged. A maximum of 126 ACL entries with deny action can be logged.
- For virtual ACL entries, the same match rule number is reused. Similarly, when an ACL entry is deleted that was previously enabled for ACL logging, the match rule number used by it is released back to the pool or available set of match indices so that it can be reused for subsequent allocations.
- If you enabled the count of packets for the ACL entry for which you configured logging, and if the logging is deactivated in a specific interval owing to the threshold having exceeded, the count of packets that exceeded the logging threshold value during that interval is logged when the subsequent log record (in the next interval) is generated for that ACL entry.
- When you delete an ACL entry, the logging settings associated with it are also removed.
- ACL logging is supported for standard and extended IPv4 ACLs, IPv6 ACLs, and standard and extended MAC ACLs.
- For ACL entries applied on port-channel interfaces, one match index for every member interface of the port-channel interface is assigned. Therefore, the total available match indices of 251 are split (125 match indices for permit action and 126 match indices for the deny action).
- You can configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs on egress interfaces.
- The total available match rule indices is 255 with four match indices used by other modules, leaving 251 indices available for ACL logging.

## Configuring ACL Logging

To configure the maximum number of ACL log messages to be generated and the frequency at which these messages must be generated, perform the following steps:



**NOTE:** This example describes the configuration of ACL logging for standard IP access lists. You can enable the logging capability for standard and extended IPv4 ACLs, IPv6 ACLs, and standard and extended MAC ACLs.

1. Specify the maximum number of ACL logs or the threshold that can be generated by using the `threshold-in-msgs count` option with the `seq`, `permit`, or `deny` commands. Upon exceeding the specified maximum limit, the generation of ACL logs is terminated. You can enter a threshold in the range of 1-100. By default, 10 ACL logs are generated if you do not specify the threshold explicitly.

CONFIG-STD-NACL mode

```
seq sequence-number {deny | permit} {source [mask] | any | host ip-address}
[log [threshold-in-msgs count]]
```

2. Specify the interval in minutes at which ACL logs must be generated. You can enter an interval in the range of 1-10 minutes. The default frequency at which ACL logs are generated is 5 minutes. If ACL logging is stopped because the configured threshold has exceeded, it is re-enabled after the logging interval period elapses. ACL logging is supported for standard and extended IPv4 ACLs, IPv6 ACLs, and standard and extended MAC ACLs. Configure ACL logging only on ACLs that are applied to ingress interfaces; you cannot enable logging for ACLs that are associated with egress interfaces.

CONFIG-STD-NACL mode

```
seq sequence-number {deny | permit} {source [mask] | any | host ip-address}
[log [interval minutes]]
```

## Flow-Based Monitoring Support for ACLs

Flow-based monitoring conserves bandwidth by monitoring only the specified traffic instead of all traffic on the interface. It is available for Layer 2 and Layer 3 ingress traffic. You can specify traffic using standard or extended access-lists. This feature copies all incoming packets on one port and forwards (mirrors) them to another port. The source port is the monitored port (MD) and the destination port is the monitoring port (MG).

The port mirroring application maintains and performs all the monitoring operations on the chassis. ACL information is sent to the ACL manager, which in turn notifies the ACL agent to add entries in the CAM area. Duplicate entries in the ACL are not saved.

When a packet arrives at a port that is being monitored, the packet is validated against the configured ACL rules. If the packet matches an ACL rule, the system examines the corresponding flow processor to perform the action specified for that port. If the mirroring action is set in the flow processor entry, the destination port details, to which the mirrored information must be sent, are sent to the destination port.

When a stack unit is reset or a stack unit undergoes a failure, the ACL agent registers with the port mirroring application. The port mirroring utility downloads the monitoring configuration to the ACL agent. The interface manager notifies the port mirroring application about the removal of an interface when an ACL entry associated with that interface to is deleted.

### Behavior of Flow-Based Monitoring

Activate flow-based monitoring for a monitoring session by entering the `flow-based enable` command in the Monitor Session mode. When you enable this capability, traffic with particular flows that

are traversing through the ingress interfaces are examined, and appropriate ACLs can be applied in the ingress direction. By default, flow-based monitoring is not enabled.

You must specify the `monitor` option with the `permit`, `deny`, or `seq` command for ACLs that are assigned to the source or the monitored port (MD) to enable the evaluation and replication of traffic that is traversing to the destination port. Enter the keyword `monitor` with the `seq`, `permit`, or `deny` command for the ACL rules to allow or drop IPv4, IPv6, ARP, UDP, EtherType, ICMP, and TCP packets. The ACL rule describes the traffic that you want to monitor, and the ACL in which you are creating the rule will be applied to the monitored interface. Flow monitoring is supported for standard and extended IPv4 ACLs, standard and extended IPv6 ACLs, and standard and extended MAC ACLs.

CONFIG-STD-NACL mode

```
seq sequence-number {deny | permit} {source [mask] | any | host ip-address}
[count [byte]] [monitor]
```

If the number of monitoring sessions increases, inter-process communication (IPC) bandwidth utilization will be high. The ACL manager might require a large bandwidth when you assign an ACL, with many entries, to an interface.

The ACL agent module saves monitoring details in its local database and also in the CAM region to monitor packets that match the specified criterion. The ACL agent maintains data on the source port, the destination port, and the endpoint to which the packet must be forwarded when a match occurs with the ACL entry.

If you configure the `flow-based enable` command and do not apply an ACL on the source port or the monitored port, both flow-based monitoring and port mirroring do not function. Flow-based monitoring is supported only for ingress traffic and not for egress packets.

The port mirroring application maintains a database that contains all monitoring sessions (including port monitor sessions). It has information regarding the sessions that are enabled for flow-based monitoring and those sessions that are not enabled for flow-based monitoring. It downloads monitoring configuration to the ACL agent whenever the ACL agent is registered with the port mirroring application or when flow-based monitoring is enabled.

The `show monitor session session-id` command has been enhanced to display the `Type` field in the output, which indicates whether a particular session is enabled for flow-monitoring.

#### Example Output of the `show` Command

```
#show running-config monitor session
!
monitor session 11
 flow-based enable
 source GigabitEthernet 13/0 destination GigabitEthernet 13/1 direction both
```

The `show running-config monitor session` command displays whether flow-based monitoring is enabled for a particular session.

The `show config` command has been modified to display monitoring configuration in a particular session.

### Example Output of the show Command

```
(conf-mon-sess-11)#show config
!
monitor session 11
 flow-based enable
source GigabitEthernet 13/0 destination GigabitEthernet 13/1 direction
both
```

The show ip | mac | ipv6 accounting commands have been enhanced to display whether monitoring is enabled for traffic that matches with the rules of the specific ACL.

### Example Output of the show Command

```
Dell# show ip accounting access-list
!
Extended Ingress IP access list kar on GigabitEthernet 10/0
Total cam count 1
seq 5 permit ip 192.168.20.0/24 173.168.20.0/24 monitor

Dell#show mac accounting access-list kar in gi 10/0 out
Egress Extended mac access-list kar on GigabitEthernet 10/0
seq 5 permit host 11:11:11:11:11:11 host 22:22:22:22:22:22 monitor
seq 10 permit host 22:22:22:22:22:22 any monitor
seq 15 permit host 00:0f:fe:1e:de:9b host 0a:0c:fb:1d:fc:aa monitor

Dell#show ipv6 accounting access-list
!
Ingress IPv6 access list kar on GigabitEthernet 10/0
Total cam count 1
seq 5 permit ipv6 22::/24 33::/24 monitor
```

## Enabling Flow-Based Monitoring

Flow-based monitoring conserves bandwidth by monitoring only specified traffic instead of all traffic on the interface. This feature is particularly useful when looking for malicious traffic. It is available for Layer 2 and Layer 3 ingress and egress traffic. You can specify traffic using standard or extended access-lists.

1. Enable flow-based monitoring for a monitoring session.

```
MONITOR SESSION mode
```

```
flow-based enable
```

2. Define access-list rules that include the keyword `monitor`. Dell Networking OS only considers port monitoring traffic that matches rules with the keyword `monitor`.

```
CONFIGURATION mode
```

```
ip access-list
```

For more information, see [Access Control Lists \(ACLs\)](#).

3. Apply the ACL to the monitored port.

```
INTERFACE mode
```

```
ip access-group access-list
```



### Example of the flow-based enable Command

To view an access-list that you applied to an interface, use the `show ip accounting access-list` command from EXEC Privilege mode.

```
Dell(conf)#monitor session 0
Dell(conf-mon-sess-0)#flow-based enable
Dell(conf)#ip access-list ext testflow
Dell(config-ext-nacl)#seq 5 permit icmp any any count bytes monitor
Dell(config-ext-nacl)#seq 10 permit ip 102.1.1.0/24 any count bytes monitor
Dell(config-ext-nacl)#seq 15 deny udp any any count bytes
Dell(config-ext-nacl)#seq 20 deny tcp any any count bytes
Dell(config-ext-nacl)#exit
Dell(conf)#interface gig 1/1
Dell(conf-if-gi-1/1)#ip access-group testflow in
Dell(conf-if-gi-1/1)#show config
!
interface GigabitEthernet 1/1
 ip address 10.11.1.254/24
 ip access-group testflow in
 shutdown
Dell(conf-if-gi-1/1)#exit
Dell(conf)#do show ip accounting access-list testflow
!
Extended Ingress IP access list testflow on GigabitEthernet 1/1
Total cam count 4
 seq 5 permit icmp any any monitor count bytes (0 packets 0 bytes)
 seq 10 permit ip 102.1.1.0/24 any monitor count bytes (0 packets 0 bytes)
 seq 15 deny udp any any count bytes (0 packets 0 bytes)
 seq 20 deny tcp any any count bytes (0 packets 0 bytes)
Dell(conf)#do show monitor session 0
SessionID Source Destination Direction Mode Type

0 Gi 1/1 Gi 1/2 rx interface Flow-based
```

# Bidirectional Forwarding Detection (BFD)

Bidirectional forwarding detection (BFD) is a protocol that is used to rapidly detect communication failures between two adjacent systems.

It is a simple and lightweight replacement for existing routing protocol link state detection mechanisms. It also provides a failure detection solution for links on which no routing protocol is used.

BFD is a simple hello mechanism. Two neighboring systems running BFD establish a session using a three-way handshake. After the session has been established, the systems exchange periodic control packets at sub-second intervals. If a system does not receive a hello packet within a specified amount of time, routing protocols are notified that the forwarding path is down.

BFD provides forwarding path failure detection times on the order of milliseconds rather than seconds as with conventional routing protocol hellos. It is independent of routing protocols, and as such, provides a consistent method of failure detection when used across a network. Networks converge faster because BFD triggers link state changes in the routing protocol sooner and more consistently because BFD eliminates the use of multiple protocol-dependent timers and methods.


BFD also carries less overhead than routing protocol hello mechanisms. Control packets can be encapsulated in any form that is convenient, and, on Dell Networking routers, BFD agents maintain sessions that reside on the line card, which frees resources on the route processor module (RPM). Only session state changes are reported to the BFD Manager (on the RPM), which in turn notifies the routing protocols that are registered with it.

BFD is an independent and generic protocol, which all media, topologies, and routing protocols can support using any encapsulation. Dell Networking has implemented BFD at Layer 3 and with user datagram protocol (UDP) encapsulation. BFD functionality will be implemented in phases. OSPF, IS-IS, VRRP, VLANs, LAGs, static routes, and physical ports support BFD, based on the IETF internet *draft draft-ietf-bfd-base-03*.

## How BFD Works

Two neighboring systems running BFD establish a session using a three-way handshake.

After the session has been established, the systems exchange control packets at agreed upon intervals. In addition, systems send a control packet anytime there is a state change or change in a session parameter. These control packets are sent without regard to transmit and receive intervals.

 **NOTE:** The Dell Networking operating system does not support multi-hop BFD sessions.

If a system does not receive a control packet within an agreed-upon amount of time, the BFD agent changes the session state to Down. It then notifies the BFD manager of the change and sends a control

packet to the neighbor that indicates the state change (though it might not be received if the link or receiving interface is faulty). The BFD manager notifies the routing protocols that are registered with it (clients) that the forwarding path is down and a link state change is triggered in all protocols.

**NOTE:** A session state change from Up to Down is the only state change that triggers a link state change in the routing protocol client.

## BFD Packet Format

Control packets are encapsulated in user datagram protocol (UDP) packets. The following illustration shows the complete encapsulation of a BFD control packet inside an IPv4 packet.

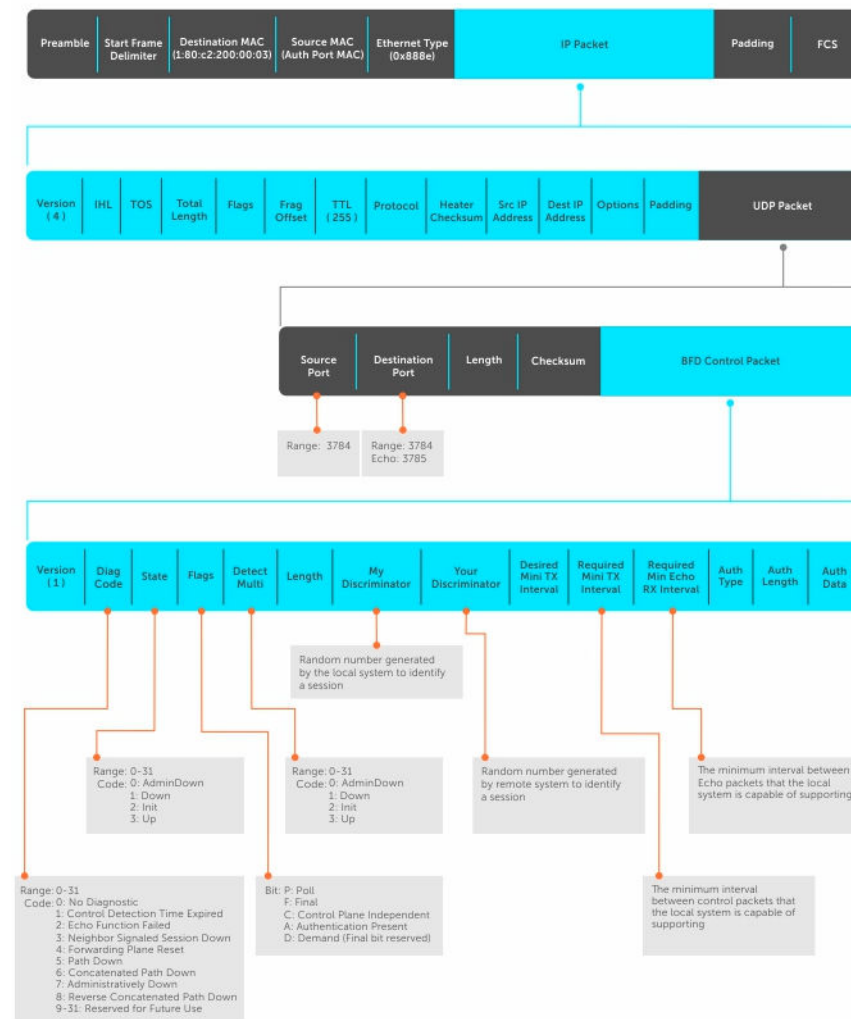





Figure 7. BFD in IPv4 Packet Format

| Field           | Description                              |
|-----------------|------------------------------------------|
| Diagnostic Code | The reason that the last session failed. |

| Field                                                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| State                                                           | The current local session state. Refer to <a href="#">BFD Sessions</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Flag                                                            | A bit that indicates packet function. If the poll bit is set, the receiving system must respond as soon as possible, without regard to its transmit interval. The responding system clears the poll bit and sets the final bit in its response. The poll and final bits are used during the handshake and in Demand mode (refer to <a href="#">BFD Sessions</a> ).<br><br> <b>NOTE:</b> The Dell Networking OS does not currently support multi-point sessions, Demand mode, authentication, or control plane independence; these bits are always clear. |
| Detection Multiplier                                            | The number of packets that must be missed in order to declare a session down.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Length                                                          | The entire length of the BFD packet.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| My Discriminator                                                | A random number generated by the local system to identify the session.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Your Discriminator                                              | A random number generated by the remote system to identify the session. Discriminator values are necessary to identify the session to which a control packet belongs because there can be many sessions running on a single interface.                                                                                                                                                                                                                                                                                                                                                                                                    |
| Desired Min TX Interval                                         | The minimum rate at which the local system would like to send control packets to the remote system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Required Min RX Interval                                        | The minimum rate at which the local system would like to receive control packets from the remote system.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Required Min Echo RX                                            | The minimum rate at which the local system would like to receive echo packets.<br><br> <b>NOTE:</b> The Dell Networking OS does not currently support the echo function.                                                                                                                                                                                                                                                                                                                                                                               |
| Authentication Type, Authentication Length, Authentication Data | An optional method for authenticating control packets.<br><br> <b>NOTE:</b> The Dell Networking OS does not currently support the BFD authentication function.                                                                                                                                                                                                                                                                                                                                                                                         |

Two important parameters are calculated using the values contained in the control packet.

|                   |                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Transmit interval | Transmit interval is the agreed-upon rate at which a system sends control packets. Each system has its own transmit interval, which is the greater of the last received remote Desired TX Interval and the local Required Min RX Interval.                                                                                                                                                                             |
| Detection time    | Detection time is the amount of time that a system does not receive a control packet, after which the system determines that the session has failed. Each system has its own detection time. <ul style="list-style-type: none"> <li>In Asynchronous mode: Detection time is the remote Detection Multiplier multiplied by greater of the remote Desired TX Interval and the local Required Min RX Interval.</li> </ul> |

- In Demand mode: Detection time is the local Detection Multiplier multiplied by the greater of the local Desired Min TX and the remote Required Min RX Interval.

## BFD Sessions

You must enable BFD on both sides of a link in order to establish a session.

The two participating systems can assume either of two roles:

|                |                                                                                                                                  |
|----------------|----------------------------------------------------------------------------------------------------------------------------------|
| <b>Active</b>  | The active system initiates the BFD session. Both systems can be active for the same session.                                    |
| <b>Passive</b> | The passive system does not initiate a session. It only responds to a request for session initialization from the active system. |

A BFD session has two modes:

|                          |                                                                                                                                                                                                                                             |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Asynchronous mode</b> | In Asynchronous mode, both systems send periodic control messages at an agreed upon interval to indicate that their session status is Up.'                                                                                                  |
| <b>Demand mode</b>       | If one system requests Demand mode, the other system stops sending periodic control packets; it only sends a response to status inquiries from the Demand mode initiator. Either system (but not both) can request Demand mode at any time. |

 **NOTE:** The Dell Networking OS supports Asynchronous mode only.

A session can have four states: Administratively Down, Down, Init, and Up.

|                              |                                                                                                                      |
|------------------------------|----------------------------------------------------------------------------------------------------------------------|
| <b>Administratively Down</b> | The local system does not participate in a particular session.                                                       |
| <b>Down</b>                  | The remote system is not sending control packets or at least not within the detection time for a particular session. |
| <b>Init</b>                  | The local system is communicating.                                                                                   |
| <b>Up</b>                    | Both systems are exchanging control packets.                                                                         |

The session is declared down if:

- A control packet is not received within the detection time.
- Sufficient echo packets are lost.
- Demand mode is active and a control packet is not received in response to a poll packet.

## BFD Three-Way Handshake

A three-way handshake must take place between the systems that participate in the BFD session.

The handshake shown in the following illustration assumes that there is one active and one passive system, and that this session is the first session established on this link. The default session state on both ports is Down.

1. The active system sends a steady stream of control packets that indicates that its session state is Down, until the passive system responds. These packets are sent at the desired transmit interval of the Active system. The Your Discriminator field is set to zero.
2. When the passive system receives any of these control packets, it changes its session state to Init and sends a response that indicates its state change. The response includes its session ID in the My Discriminator field and the session ID of the remote system in the Your Discriminator field.
3. The active system receives the response from the passive system and changes its session state to Up. It then sends a control packet indicating this state change. This is the third and final part of the handshake. Now the discriminator values have been exchanged and the transmit intervals have been negotiated.
4. The passive system receives the control packet and changes its state to Up. Both systems agree that a session has been established. However, because both members must send a control packet — that requires a response — anytime there is a state change or change in a session parameter, the passive system sends a final response indicating the state change. After this, periodic control packets are exchanged.

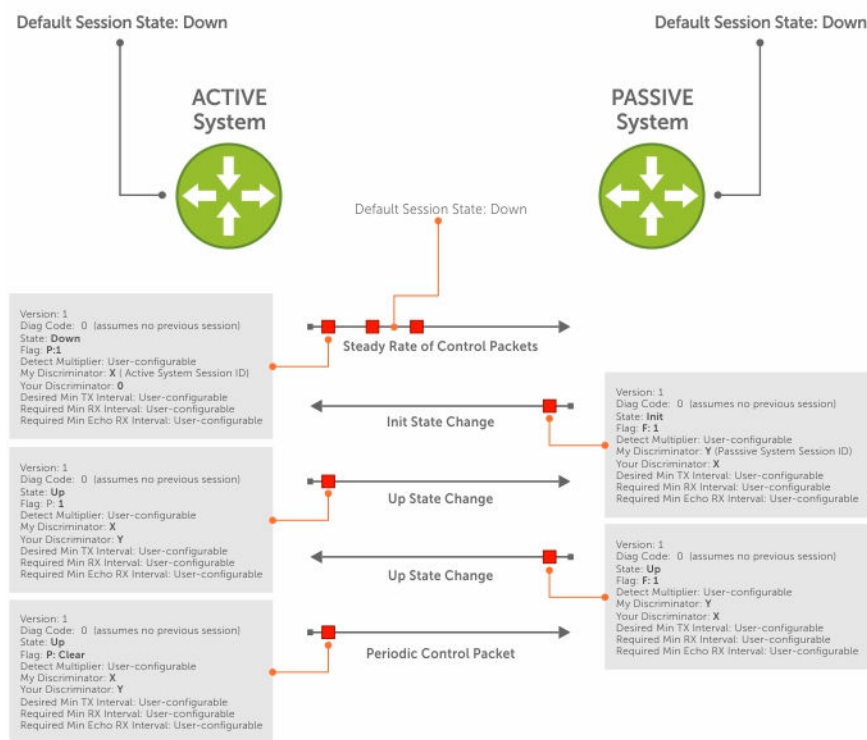


Figure 8. BFD Three-Way Handshake State Changes

## Session State Changes

The following illustration shows how the session state on a system changes based on the status notification it receives from the remote system. For example, if a session on a system is down and it

receives a Down status notification from the remote system, the session state on the local system changes to Init.




Figure 9. Session State Changes

## Important Points to Remember

- BFD for line card ports is hitless, but is not hitless for VLANs because they are instantiated on the RPM.
- The Dell Networking OS supports a maximum of 100 sessions per BFD agent. Each linecard processor has a BFD Agent, so the limit translates to 100 BFD sessions per linecard.
- Enable BFD on both ends of a link.
- Demand mode, authentication, and the Echo function are not supported.
- BFD is not supported on multi-hop and virtual links.
- Protocol Liveness is supported for routing protocols only.
- The Dell Networking OS supports only OSPF, OSPFv3, BGP, and VRRP protocols as BFD clients.

# Configure BFD

This section contains the following procedures.

- [Configure BFD for Physical Ports](#)
- [Configure BFD for Port-Channels](#)
- [Configure BFD for Static Routes](#)
- [Configure BFD for OSPF](#)
- [Configure BFD for OSPFv3](#)
- [Configure BFD for BGP](#)
- [Configure BFD for VRRP](#)
- [Configure BFD for VLANs](#)
- [Configuring Protocol Liveness](#)
- [Troubleshooting BFD](#)

## Configure BFD for Physical Ports

BFD on physical ports is useful when you do not enable the routing protocol.

Without BFD, if the remote system fails, the local system does not remove the connected route until the first failed attempt to send a packet. When you enable BFD, the local system removes the route as soon as it stops receiving periodic control packets from the remote system.

Configuring BFD for a physical port is a two-step process:

1. Enable BFD globally. Refer to [Enabling BFD Globally](#).
2. Establish a session with a next-hop neighbor.

### Related Configuration Tasks

- [Changing Physical Port Session Parameters](#).
- [Disabling and Re-Enabling BFD](#).

## Enabling BFD Globally

You must enable BFD globally on both routers.

To enable the BFD globally, use the following command.

- Enable BFD globally.  
CONFIGURATION mode

```
bfd enable
```

### Example of Verifying BFD is Enabled

To verify that BFD is enabled globally, use the `show running bfd` command.



The bold line shows that BFD is enabled.

```
R1(conf)#bfd ?
enable Enable BFD protocol
protocol-liveness Enable BFD protocol-liveness
R1(conf)#bfd enable

R1(conf)#do show running-config bfd
!
bfd enable
R1(conf)#
```

## Establishing a Session on Physical Ports

To establish a session, enable BFD at the interface level on both ends of the link, as shown in the following illustration. The configuration parameters do not need to match.

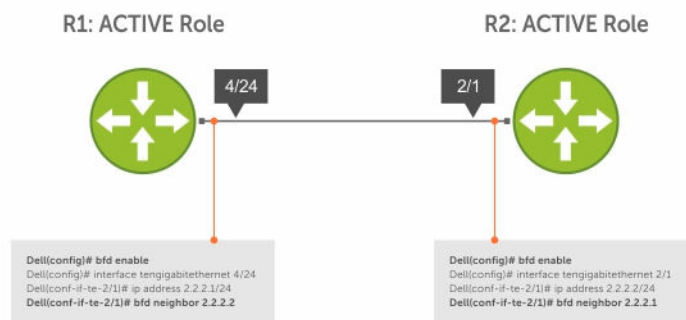


Figure 10. Establishing a BFD Session on Physical Ports

1. Enter interface mode.  
CONFIGURATION mode  

```
interface
```
2. Assign an IP address to the interface if one is not already assigned.  
INTERFACE mode  

```
ip address ip-address
```
3. Identify the neighbor that the interface participates with the BFD session.  
INTERFACE mode  

```
bfd neighbor ip-address
```

### Example of Verifying Session Creation

#### Example of the Viewing Detailed BFD Session Information

To verify that the session is established, use the `show bfd neighbors` command.

The bold line shows the BFD session.

```
R1(conf-if-gi-4/24)#do show bfd neighbors
* - Active session role
Ad Dn - Admin Down
```

```

C - CLI
I - ISIS
O - OSPF
R - Static Route (RTM)
LocalAddr RemoteAddr Interface State Rx-int Tx-int Mult Clients
* 2.2.2.1 2.2.2.2 Gi 4/24 Up 100 100 3 C

```

To view specific information about BFD sessions, use the `show bfd neighbors detail` command.

```
R1(conf-if-gi-4/24)#do show bfd neighbors detail
```

```

Session Discriminator: 1
Neighbor Discriminator: 1
Local Addr: 2.2.2.1
Local MAC Addr: 00:01:e8:09:c3:e5
Remote Addr: 2.2.2.2
Remote MAC Addr: 00:01:e8:06:95:a2
Int: GigabitEthernet 4/24
State: Up
Configured parameters:
 TX: 100ms, RX: 100ms, Multiplier: 3
Neighbor parameters:
 TX: 100ms, RX: 100ms, Multiplier: 3
Actual parameters:
 TX: 100ms, RX: 100ms, Multiplier: 3
Role: Active
Delete session on Down: False
Client Registered: CLI
Uptime: 00:03:57
Statistics:
 Number of packets received from neighbor: 1775
 Number of packets sent to neighbor: 1775
 Number of state changes: 1
 Number of messages from IFA about port state change: 0
 Number of messages communicated b/w Manager and Agent: 4

```

Log messages display when you configure both interfaces for BFD.

```


R1(conf-if-gi-4/24)#00:36:01: %RPM0-P:RP2 %BFDMGR-1-BFD_STATE_CHANGE:
Changed session state to
Down for neighbor 2.2.2.2 on interface Gi 4/24 (diag: 0)
00:36:02: %RPM0-P:RP2 %BFDMGR-1-BFD_STATE_CHANGE: Changed session state to
Up for neighbor
2.2.2.2 on interface Gi 4/24 (diag: 0)

```

## Changing Physical Port Session Parameters

Configure BFD sessions with default intervals and a default role (active).

The parameters that you can configure are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. Configure these parameters per interface; if you change a parameter, the change affects all physical port sessions on that interface.

 **NOTE:** Dell Networking recommends maintaining the default values.

Change session parameters for all sessions on an interface.

INTERFACE mode

```
bfd interval milliseconds min_rx milliseconds multiplier value role [active
| passive]
```

### Changing Session Parameters for Physical Ports

View session parameters using the `show bfd neighbors detail` command.

The bold line shows the parameter changes.

```
R1(conf-if-gi-4/24)#bfd interval 100 min_rx 100 multiplier 4 role passive
R1(conf-if-gi-4/24)#do show bfd neighbors detail
```

```
Session Discriminator: 1
Neighbor Discriminator: 1
Local Addr: 2.2.2.1
Local MAC Addr: 00:01:e8:09:c3:e5
Remote Addr: 2.2.2.2
Remote MAC Addr: 00:01:e8:06:95:a2
Int: GigabitEthernet 4/24
State: Up
Configured parameters:
TX: 100ms, RX: 100ms, Multiplier: 4
Neighbor parameters:
TX: 100ms, RX: 100ms, Multiplier: 3
Actual parameters:
TX: 100ms, RX: 100ms, Multiplier: 4
Role: Passive
Delete session on Down: False
Client Registered: CLI
Uptime: 00:09:06
Statistics:
Number of packets received from neighbor: 4092
Number of packets sent to neighbor: 4093
Number of state changes: 1
Number of messages from IFA about port state change: 0
Number of messages communicated b/w Manager and Agent: 7
```

### Disabling and Re-Enabling BFD

BFD is enabled on all interfaces by default, though sessions are not created unless explicitly configured. If you disable BFD, all of the sessions on that interface are placed in an Administratively Down state ( the first message example), and the remote systems are notified of the session state change (the second message example).

To disable and re-enable BFD on an interface, use the following commands.

- Disable BFD on an interface.

```
INTERFACE mode
```

```
no bfd enable
```

- Enable BFD on an interface.

```
INTERFACE mode
```

```
bfd enable
```

If you disable BFD on a local interface, this message displays:

```
R1(conf-if-gi-4/24)#01:00:52: %RPM0-P:RP2 %BFDMGR-1-BFD_STATE_CHANGE:
Changed session state to Ad
Dn for neighbor 2.2.2.2 on interface Gi 4/24 (diag: 0)
```

If the remote system state changes due to the local state administration being down, this message displays:

```
R2>01:32:53: %RPM0-P:RP2 %BFDMGR-1-BFD_STATE_CHANGE: Changed session state to Down for neighbor 2.2.2.1 on interface Gi 2/1 (diag: 7)
```

## Configure BFD for Static Routes

BFD offers systems a link state detection mechanism for static routes.

With BFD, systems are notified to remove static routes from the routing table as soon as the link state change occurs, rather than waiting until packets fail to reach their next hop.

Configuring BFD for static routes is a three-step process:

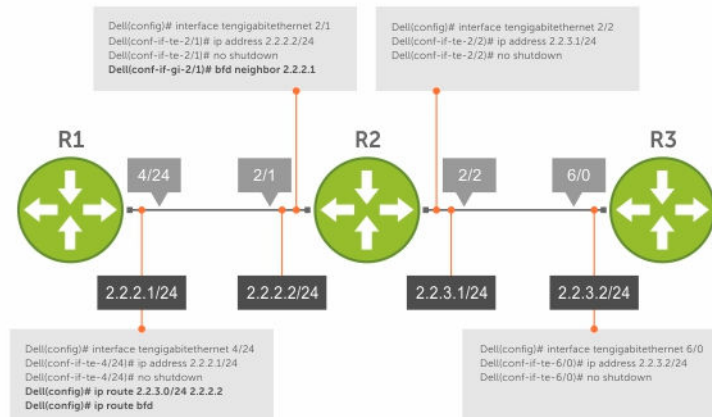
1. Enable BFD globally. Refer to [Enabling BFD Globally](#).
2. On the local system, establish a session with the next hop of a static route. Refer to [Establishing Sessions for Static Routes](#).
3. On the remote system, establish a session with the physical port that is the origin of the static route. Refer to [Establishing a Session on Physical Ports](#).

## Related Configuration Tasks

- [Changing Static Route Session Parameters](#)
- [Disabling BFD for Static Routes](#)

## Establishing Sessions for Static Routes

Sessions are established for all neighbors that are the next hop of a static route.



**Figure 11. Establishing Sessions for Static Routes**

To establish a BFD session, use the following command.

- Establish BFD sessions for all neighbors that are the next hop of a static route.  
CONFIGURATION mode

```
ip route bfd
```

### Example of the show bfd neighbors Command to Verify Static Routes

To verify that sessions have been created for static routes, use the `show bfd neighbors` command.

The bold line shows BFD for static routes is enabled.

```
R1(conf)#ip route 2.2.3.0/24 2.2.2.2
R1(conf)#ip route bfd
R1(conf)#do show bfd neighbors

* - Active session role
Ad Dn - Admin Down
C - CLI
I - ISIS
O - OSPF
R - Static Route (RTM)
LocalAddr RemoteAddr Interface State Rx-int Tx-int Mult Clients
2.2.2.1 2.2.2.2 Gi 4/24 Up 100 100 4 R
```

To view detailed session information, use the `show bfd neighbors detail` command, as shown in the examples in [Disabling BFD for BGP](#).

## Changing Static Route Session Parameters

BFD sessions are configured with default intervals and a default role.

The parameters you can configure are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. These parameters are configured for all static routes. If you change a parameter, the change affects all sessions for static routes.

To change parameters for static route sessions, use the following command .

- Change parameters for all static route sessions.

CONFIGURATION mode

```
ip route bfd interval milliseconds min_rx milliseconds multiplier value role
[active | passive]
```

To view session parameters, use the `show bfd neighbors detail` command, as shown in the examples in [Displaying BFD for BGP Information](#).

## Disabling BFD for Static Routes

If you disable BFD, all static route BFD sessions are torn down.

A final Admin Down packet is sent to all neighbors on the remote systems, and those neighbors change to the Down state.

To disable BFD for static routes, use the following command.

- Disable BFD for static routes.

CONFIGURATION mode

```
no ip route bfd
```

## Configure BFD for OSPF

When using BFD with OSPF, the OSPF protocol registers with the BFD manager on the RPM.

BFD sessions are established with all neighboring interfaces participating in OSPF. If a neighboring interface fails, the BFD agent on the line card notifies the BFD manager, which in turn notifies the OSPF protocol that a link state change occurred.

Configuring BFD for OSPF is a two-step process:

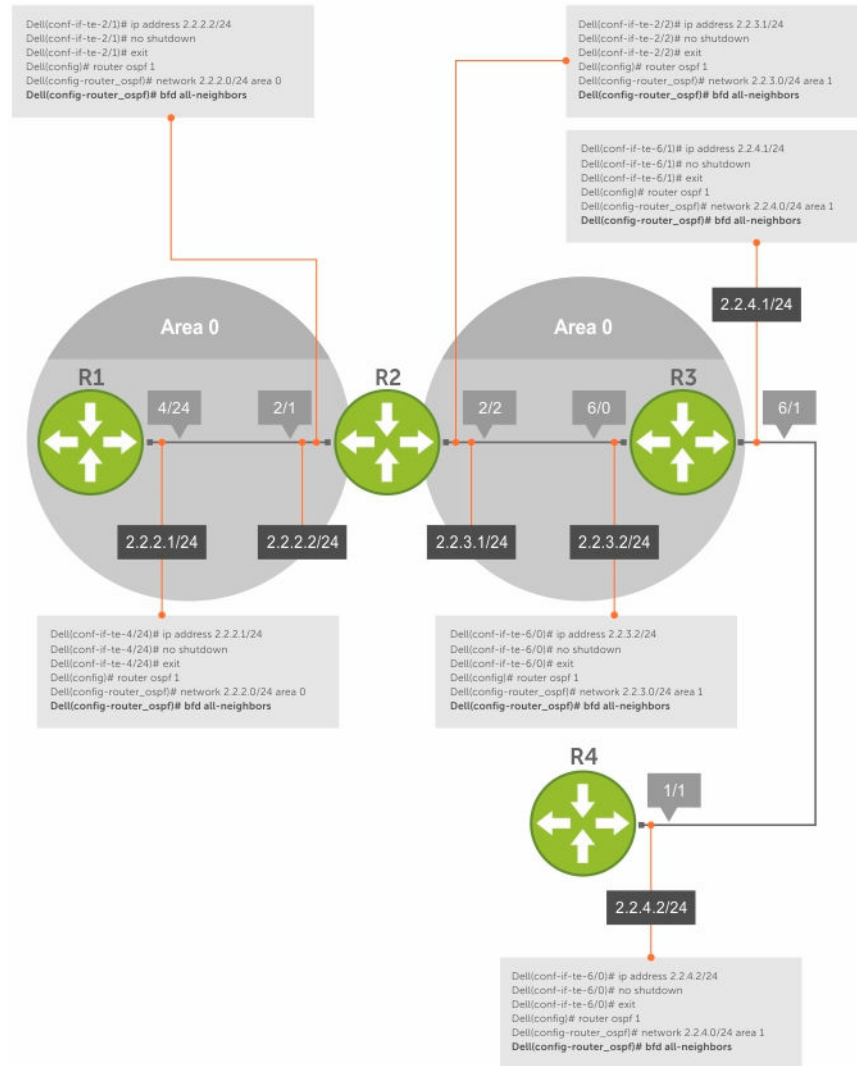
1. Enable BFD globally. Refer to [Enabling BFD Globally](#).
2. Establish sessions with OSPF neighbors. Refer to [Establishing Sessions with OSPF Neighbors](#).

### Related Configuration Tasks

- [Changing OSPF Session Parameters](#)
- [Disabling BFD for OSPF](#)

## Establishing Sessions with OSPF Neighbors

BFD sessions can be established with all OSPF neighbors at once or sessions can be established with all neighbors out of a specific interface. Sessions are only established when the OSPF adjacency is in the Full state.



**Figure 12. Establishing Sessions with OSPF Neighbors**

To establish BFD with all OSPF neighbors or with OSPF neighbors on a single interface, use the following commands.

- Establish sessions with all OSPF neighbors.  
ROUTER-OSPF mode  
  
bfd all-neighbors
- Establish sessions with OSPF neighbors on a single interface.

INTERFACE mode

```
ip ospf bfd all-neighbors
```

### Example of Verifying Sessions with OSPF Neighbors

To view the established sessions, use the `show bfd neighbors` command.

The bold line shows the OSPF BFD sessions.

```
R2(conf-router_ospf)#bfd all-neighbors
R2(conf-router_ospf)#do show bfd neighbors

* - Active session role
Ad Dn - Admin Down
C - CLI
I - ISIS
O - OSPF
R - Static Route (RTM)

LocalAddr RemoteAddr Interface State Rx-int Tx-int Mult Clients
* 2.2.2.2 2.2.2.1 Gi 2/1 Up 100 100 3 O
* 2.2.3.1 2.2.3.2 Gi 2/2 Up 100 100 3 O
```

### Changing OSPF Session Parameters

Configure BFD sessions with default intervals and a default role.

The parameters that you can configure are: `desired tx interval`, `required min rx interval`, `detection multiplier`, and `system role`. Configure these parameters for all OSPF sessions or all OSPF sessions on a particular interface. If you change a parameter globally, the change affects all OSPF neighbors sessions. If you change a parameter at the interface level, the change affects all OSPF sessions on that interface.

To change parameters for all OSPF sessions or for OSPF sessions on a single interface, use the following commands.

- Change parameters for OSPF sessions.

ROUTER-OSPF mode

```
bfd all-neighbors interval milliseconds min_rx milliseconds multiplier value
role [active | passive]
```

- Change parameters for all OSPF sessions on an interface.

INTERFACE mode

```
ip ospf bfd all-neighbors interval milliseconds min_rx milliseconds
multiplier value role [active | passive]
```

To view session parameters, use the `show bfd neighbors detail` command, as shown in the example in [Displaying BFD for BGP Information](#).



## Disabling BFD for OSPF

If you disable BFD globally, all sessions are torn down and sessions on the remote system are placed in a Down state.

If you disable BFD on an interface, sessions on the interface are torn down and sessions on the remote system are placed in a Down state. Disabling BFD does not trigger a change in BFD clients; a final Admin Down packet is sent before the session is terminated.

To disable BFD sessions, use the following commands.

- Disable BFD sessions with all OSPF neighbors.  
ROUTER-OSPF mode  

```
no bfd all-neighbors
```
- Disable BFD sessions with all OSPF neighbors on an interface.  
INTERFACE mode  

```
ip ospf bfd all-neighbors disable
```

## Configure BFD for OSPFv3

BFD for OSPFv3 provides support for IPV6.

Configuring BFD for OSPFv3 is a two-step process:

1. Enable BFD globally.
2. Establish sessions with OSPFv3 neighbors.

### Related Configuration Tasks

- [Changing OSPFv3 Session Parameters](#)
- [Disabling BFD for OSPFv3](#)

## Establishing Sessions with OSPFv3 Neighbors

You can establish BFD sessions with all OSPFv3 neighbors at once or with all neighbors out of a specific interface. Sessions are only established when the OSPFv3 adjacency is in the Full state.

To establish BFD with all OSPFv3 neighbors or with OSPFv3 neighbors on a single interface, use the following commands.

- Establish sessions with all OSPFv3 neighbors.  
ROUTER-OSPFv3 mode  

```
bfd all-neighbors
```
- Establish sessions with OSPFv3 neighbors on a single interface.  
INTERFACE mode  

```
ipv6 ospf bfd all-neighbors
```

To view the established sessions, use the `show bfd neighbors` command.

## Changing OSPFv3 Session Parameters

Configure BFD sessions with default intervals and a default role.

The parameters that you can configure are: `desired tx interval`, `required min rx interval`, `detection multiplier`, and `system role`. Configure these parameters for all OSPFv3 sessions or all OSPFv3 sessions on a particular interface. If you change a parameter globally, the change affects all OSPFv3 neighbors sessions. If you change a parameter at the interface level, the change affects all OSPFv3 sessions on that interface.

To change parameters for all OSPFv3 sessions or for OSPFv3 sessions on a single interface, use the following commands.

To view session parameters, use the `show bfd neighbors detail` command, as shown in the example in [Displaying BFD for BGP Information](#).

- Change parameters for all OSPFv3 sessions.

ROUTER-OSPFv3 mode

```
bfd all-neighbors interval milliseconds min_rx milliseconds multiplier value
role [active | passive]
```

- Change parameters for OSPFv3 sessions on a single interface.

INTERFACE mode

```
ipv6 ospf bfd all-neighbors interval milliseconds min_rx milliseconds
multiplier value role [active | passive]
```

## Disabling BFD for OSPFv3

If you disable BFD globally, all sessions are torn down and sessions on the remote system are placed in a Down state.

If you disable BFD on an interface, sessions on the interface are torn down and sessions on the remote system are placed in a Down state. Disabling BFD does not trigger a change in BFD clients; a final Admin Down packet is sent before the session is terminated.

To disable BFD sessions, use the following commands.

- Disable BFD sessions with all OSPFv3 neighbors.

ROUTER-OSPFv3 mode

```
no bfd all-neighbors
```

- Disable BFD sessions with OSPFv3 neighbors on a single interface.

INTERFACE mode

```
ipv6 ospf bfd all-neighbors disable
```

## Configure BFD for BGP

In a BGP core network, bidirectional forwarding detection (BFD) provides rapid detection of communication failures in BGP fast-forwarding paths between internal BGP (iBGP) and external BGP (eBGP) peers for faster network reconvergence.

BFD for BGP is supported on 1GE, 10GE, 40GE, port-channel, and VLAN interfaces. BFD for BGP does not support IPv6 and the BGP multihop feature.

### Prerequisites

Before configuring BFD for BGP, you must first configure the following settings:

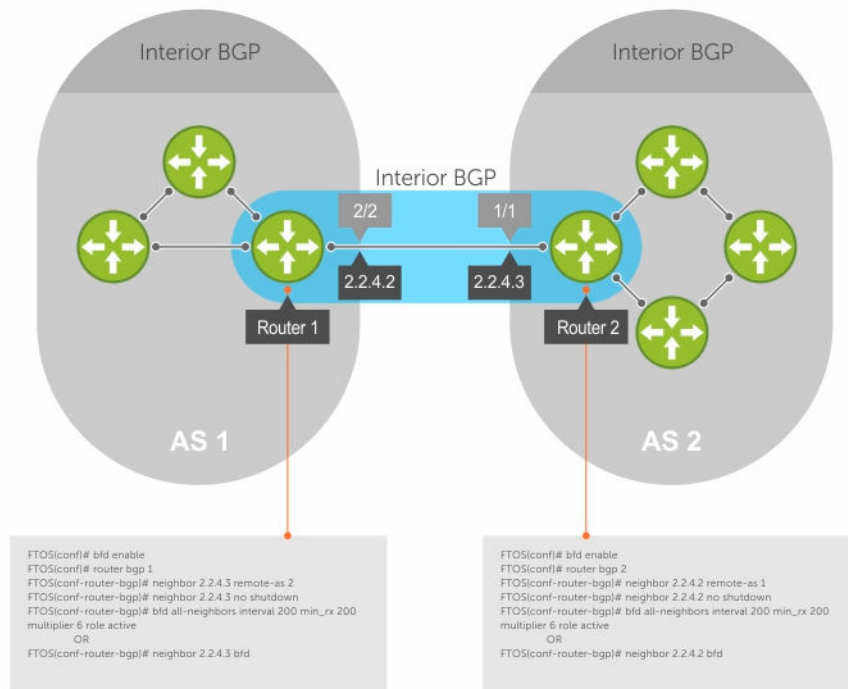
1. Configure BGP on the routers that you want to interconnect, as described in [Border Gateway Protocol IPv4 \(BGPv4\)](#).
2. Enable fast fall-over for BGP neighbors to reduce convergence time (the `neighbor fall-over` command), as described in [Configuring BGP Fast Fail-Over](#).

### Establishing Sessions with BGP Neighbors

Before configuring BFD for BGP, you must first configure BGP on the routers that you want to interconnect.

For more information, refer to [Border Gateway Protocol IPv4 \(BGPv4\)](#).

For example, the following illustration shows a sample BFD configuration on Router 1 and Router 2 that use eBGP in a transit network to interconnect AS1 and AS2. The eBGP routers exchange information with each other as well as with iBGP routers to maintain connectivity and accessibility within each autonomous system.



**Figure 13. Establishing Sessions with BGP Neighbors**

The sample configuration shows alternative ways to establish a BFD session with a BGP neighbor:

- By establishing BFD sessions with all neighbors discovered by BGP (the `bfd all-neighbors` command).
- By establishing a BFD session with a specified BGP neighbor (the `neighbor {ip-address | peer-group-name} bfd` command)

BFD packets originating from a router are assigned to the highest priority egress queue to minimize transmission delays. Incoming BFD control packets received from the BGP neighbor are assigned to the highest priority queue within the control plane policing (CoPP) framework to avoid BFD packets drops due to queue congestion.

BFD notifies BGP of any failure conditions that it detects on the link. Recovery actions are initiated by BGP.

BFD for BGP is supported only on directly-connected BGP neighbors and only in BGP IPv4 networks.

As long as each BFD for BGP neighbor receives a BFD control packet within the configured BFD interval for failure detection, the BFD session remains up and BGP maintains its adjacencies. If a BFD for BGP neighbor does not receive a control packet within the detection interval, the router informs any clients of the BFD session (other routing protocols) about the failure. It then depends on the individual routing protocols that uses the BGP link to determine the appropriate response to the failure condition. The typical response is to terminate the peering session for the routing protocol and reconverge by bypassing the failed neighboring router. A log message is generated whenever BFD detects a failure condition.

You can configure BFD for BGP on the following types of interfaces: physical port (10GE or 40GE), port channel, and VLAN.

1. Enable BFD globally.

CONFIGURATION mode

```
bfd enable
```

2. Specify the AS number and enter ROUTER BGP configuration mode.

CONFIGURATION mode

```
router bgp as-number
```

3. Add a BGP neighbor or peer group in a remote AS.

CONFIG-ROU TERBGP mode

```
neighbor {ip-address | peer-group name} remote-as as-number
```

4. Enable the BGP neighbor.

CONFIG-ROU TERBGP mode

```
neighbor {ip-address | peer-group-name} no shutdown
```

5. Configure parameters for a BFD session established with all neighbors discovered by BGP. OR Establish a BFD session with a specified BGP neighbor or peer group using the default BFD session parameters.

CONFIG-ROU TERBGP mode

```
bfd all-neighbors [interval millisecs min_rx millisecs multiplier value role
{active | passive}]
```

OR

```
neighbor {ip-address | peer-group-name} bfd
```

NOTES:

- When you establish a BFD session with a specified BGP neighbor or peer group using the `neighbor bfd` command, the default BFD session parameters are used (interval: 100 milliseconds, min\_rx: 100 milliseconds, multiplier: 3 packets, and role: active).
  - When you explicitly enable or disable a BGP neighbor for a BFD session with the `neighbor bfd` or `neighbor bfd disable` commands, the neighbor does not inherit the BFD enable/disable values configured with the `bfd all-neighbors` command or configured for the peer group to which the neighbor belongs. Also, the neighbor only inherits the global timer values configured with the `bfd all-neighbors` command (interval, min\_rx, and multiplier).
6. Repeat Steps 1 to 5 on each BGP peer participating in a BFD session.

## Disabling BFD for BGP

You can disable BFD for BGP.

To disable a BFD for BGP session with a specified neighbor, use the first command. To remove the disabled state of a BFD for BGP session with a specified neighbor, use the `no neighbor {ip-address | peer-group-name} bfd disable` command in ROUTER BGP configuration mode.

The BGP link with the neighbor returns to normal operation and uses the BFD session parameters globally configured with the `bfd all-neighbors` command or configured for the peer group to which the neighbor belongs.

- Disable a BFD for BGP session with a specified neighbor.

```
ROUTER BGP mode
```

```
neighbor {ip-address | peer-group-name} bfd disable
```

- Remove the disabled state of a BFD for BGP session with a specified neighbor.

```
ROUTER BGP mode
```

```
no neighbor {ip-address | peer-group-name} bfd disable
```

## Use BFD in a BGP Peer Group

You can establish a BFD session for the members of a peer group (the `neighbor peer-group-name bfd` command in ROUTER BGP configuration mode).

Members of the peer group may have BFD:

- Explicitly enabled (the `neighbor ip-address bfd` command)
- Explicitly disabled (the `neighbor ip-address bfd disable` command)
- Inherited (neither explicitly enabled or disabled) according to the current BFD configuration of the peer group. For information about BGP peer groups, refer to [Configuring Peer Groups](#).

If you explicitly enable (or disable) a BGP neighbor for BFD that belongs to a peer group:

- The neighbor does not inherit the BFD enable/disable values configured with the `bfd all-neighbors` command or configured for the peer group to which the neighbor belongs.
- The neighbor inherits only the global timer values that are configured with the `bfd all-neighbors` command (interval, min\_rx, and multiplier).

If you explicitly enable (or disable) a peer group for BFD that has no BFD parameters configured (for example, advertisement interval) using the `neighbor peer-group-name bfd` command, the peer group inherits any BFD settings configured with the `bfd all-neighbors` command.

## Displaying BFD for BGP Information

You can display related information for BFD for BGP.

To display information about BFD for BGP sessions on a router, use the following commands and refer to the following examples.

- Verify a BFD for BGP configuration.

```
EXEC Privilege mode
```

```
show running-config bgp
```

- Verify that a BFD for BGP session has been successfully established with a BGP neighbor. A line-by-line listing of established BFD adjacencies is displayed.

```
EXEC Privilege mode
```

```
show bfd neighbors [interface] [detail]
```

- Display BFD packet counters for sessions with BGP neighbors.  
EXEC Privilege mode  
  
show bfd counters bgp [interface]
- Check to see if BFD is enabled for BGP connections.  
EXEC Privilege mode  
  
show ip bgp summary
- Displays routing information exchanged with BGP neighbors, including BFD for BGP sessions.  
EXEC Privilege mode  
  
show ip bgp neighbors [ip-address]

### Example of Verifying BGP Configuration

### Example of Viewing All BFD Neighbors

### Example of Viewing BFD Neighbor Detail

### Example of Viewing Configured BFD Counters

### Example of Viewing BFD Summary Information

### Example of Viewing BFD Information for a Specified Neighbor

```
R2# show running-config bgp
!
router bgp 2
 neighbor 1.1.1.2 remote-as 1
 neighbor 1.1.1.2 no shutdown
 neighbor 2.2.2.2 remote-as 1
 neighbor 2.2.2.2 no shutdown
 neighbor 3.3.3.2 remote-as 1
 neighbor 3.3.3.2 no shutdown
 bfd all-neighbors
```

```
R2# show bfd neighbors
```

```
* - Active session role
Ad Dn - Admin Down
B - BGP
C - CLI
I - ISIS
O - OSPF
R - Static Route (RTM)
M - MPLS
V - VRRP
```

```
LocalAddr RemoteAddr Interface State Rx-int Tx-int Mult Clients
* 1.1.1.3 1.1.1.2 Te 6/0 Up 100 100 3 B
* 2.2.2.3 2.2.2.2 Te 6/1 Up 100 100 3 B
* 3.3.3.3 3.3.3.2 Te 6/2 Up 100 100 3 B
```

The bold lines show the BFD session parameters: TX (packet transmission), RX (packet reception), and multiplier (maximum number of missed packets).

```
R2# show bfd neighbors detail
```

```
Session Discriminator: 9
```

Neighbor Discriminator: 10  
Local Addr: 1.1.1.3  
Local MAC Addr: 00:01:e8:66:da:33  
Remote Addr: 1.1.1.2  
Remote MAC Addr: 00:01:e8:8a:da:7b  
Int: TenGigabitEthernet 6/0  
State: Up  
Configured parameters:  
**TX: 100ms, RX: 100ms, Multiplier: 3**  
Neighbor parameters:  
**TX: 100ms, RX: 100ms, Multiplier: 3**  
Actual parameters:  
**TX: 100ms, RX: 100ms, Multiplier: 3**  
Role: Active  
Delete session on Down: True  
Client Registered: BGP  
Uptime: 00:07:55  
Statistics:  
Number of packets received from neighbor: 4762  
Number of packets sent to neighbor: 4490  
Number of state changes: 2  
Number of messages from IFA about port state change: 0  
Number of messages communicated b/w Manager and Agent: 5

Session Discriminator: 10  
Neighbor Discriminator: 11  
Local Addr: 2.2.2.3  
Local MAC Addr: 00:01:e8:66:da:34  
Remote Addr: 2.2.2.2  
Remote MAC Addr: 00:01:e8:8a:da:7b  
Int: TenGigabitEthernet 6/1  
State: Up  
Configured parameters:  
TX: 100ms, RX: 100ms, Multiplier: 3  
Neighbor parameters:  
TX: 100ms, RX: 100ms, Multiplier: 3  
Actual parameters:  
TX: 100ms, RX: 100ms, Multiplier: 3  
Role: Active  
Delete session on Down: True  
Client Registered: BGP  
Uptime: 00:02:22  
Statistics:  
Number of packets received from neighbor: 1428  
Number of packets sent to neighbor: 1428  
Number of state changes: 1  
Number of messages from IFA about port state change: 0  
Number of messages communicated b/w Manager and Agent: 4

R2# show bfd counters bgp

Interface TenGigabitEthernet 6/0

Protocol BGP  
Messages:  
Registration : 5  
De-registration : 4  
Init : 0  
Up : 6  
Down : 0  
Admin Down : 2

Interface TenGigabitEthernet 6/1



```

Protocol BGP
Messages:
Registration : 5
De-registration : 4
Init : 0
Up : 6
Down : 0
Admin Down : 2

```

Interface TenGigabitEthernet 6/2

```

Protocol BGP
Messages:
Registration : 1
De-registration : 0
Init : 0
Up : 1
Down : 0
Admin Down : 2

```

The bold line shows the message displayed when you enable BFD for BGP connections.

```

R2# show ip bgp summary
BGP router identifier 10.0.0.1, local AS number 2
BGP table version is 0, main routing table version 0
BFD is enabled, Interval 100 Min_rx 100 Multiplier 3 Role Active
3 neighbor(s) using 24168 bytes of memory

```

| Neighbor | AS | MsgRcvd | MsgSent | TblVer | InQ | OutQ | Up/Down  | State/Pfx |
|----------|----|---------|---------|--------|-----|------|----------|-----------|
| 1.1.1.2  | 1  | 282     | 281     | 0      | 0   | 0    | 00:38:12 | 0         |
| 2.2.2.2  | 1  | 273     | 273     | 0      | 0   | (0)  | 04:32:26 | 0         |
| 3.3.3.2  | 1  | 282     | 281     | 0      | 0   | 0    | 00:38:12 | 0         |

The bold lines show the message displayed when you enable a BFD session with different configurations:

- Message displayed when you enable a BFD session with a BGP neighbor that inherits the global BFD session settings configured with the `global bfd all-neighbors` command.
- Message displayed when you enable a BFD session with a BGP neighbor using the `neighbor ip-address bfd` command.
- Message displayed when you enable a BGP neighbor in a peer group for which you enabled a BFD session using the `neighbor peer-group-name bfd` command

```
R2# show ip bgp neighbors 2.2.2.2
```

```

BGP neighbor is 2.2.2.2, remote AS 1, external link
BGP version 4, remote router ID 12.0.0.4
BGP state ESTABLISHED, in this state for 00:05:33
Last read 00:00:30, last write 00:00:30
Hold time is 180, keepalive interval is 60 seconds
Received 8 messages, 0 in queue
 1 opens, 0 notifications, 0 updates
 7 keepalives, 0 route refresh requests
Sent 9 messages, 0 in queue
 2 opens, 0 notifications, 0 updates
 7 keepalives, 0 route refresh requests
Minimum time between advertisement runs is 30 seconds
Minimum time before advertisements start is 0 seconds

```

```
Capabilities received from neighbor for IPv4 Unicast :
```

```

MULTIPROTO_EXT(1)
ROUTE_REFRESH(2)
CISCO_ROUTE_REFRESH(128)

Capabilities advertised to neighbor for IPv4 Unicast :
MULTIPROTO_EXT(1)
ROUTE_REFRESH(2)
CISCO_ROUTE_REFRESH(128)

Neighbor is using BGP global mode BFD configuration

For address family: IPv4 Unicast
BGP table version 0, neighbor version 0
Prefixes accepted 0 (consume 0 bytes), withdrawn 0 by peer, martian prefixes
ignored 0
Prefixes advertised 0, denied 0, withdrawn 0 from peer

Connections established 1; dropped 0
Last reset never
Local host: 2.2.2.3, Local port: 63805
Foreign host: 2.2.2.2, Foreign port: 179
E1200i_ExaScale#

R2# show ip bgp neighbors 2.2.2.3

BGP neighbor is 2.2.2.3, remote AS 1, external link
Member of peer-group pgl for session parameters
BGP version 4, remote router ID 12.0.0.4
BGP state ESTABLISHED, in this state for 00:05:33
...
Neighbor is using BGP neighbor mode BFD configuration
Peer active in peer-group outbound optimization
...

R2# show ip bgp neighbors 2.2.2.4

BGP neighbor is 2.2.2.4, remote AS 1, external link
Member of peer-group pgl for session parameters
BGP version 4, remote router ID 12.0.0.4
BGP state ESTABLISHED, in this state for 00:05:33
...
Neighbor is using BGP peer-group mode BFD configuration
Peer active in peer-group outbound optimization
...

```

## Configure BFD for VRRP

When using BFD with VRRP, the VRRP protocol registers with the BFD manager on the route processor module (RPM).

BFD sessions are established with all neighboring interfaces participating in VRRP. If a neighboring interface fails, the BFD agent on the line card notifies the BFD manager, which in turn notifies the VRRP protocol that a link state change occurred.

Configuring BFD for VRRP is a three-step process:

1. Enable BFD globally. Refer to [Enabling BFD Globally](#).
2. Establish VRRP BFD sessions with all VRRP-participating neighbors. Refer to [Establishing VRRP Sessions on VRRP Neighbors](#).

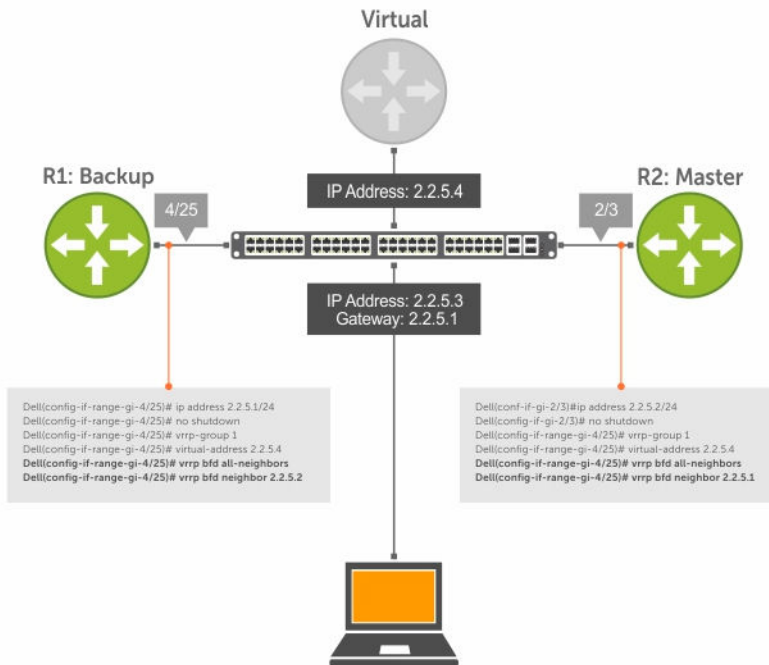
- On the master router, establish a VRRP BFD sessions with the backup routers. Refer to [Establishing Sessions with All VRRP Neighbors](#).

## Related Configuration Tasks

- [Changing VRRP Session Parameters](#).
- [Disabling BFD for VRRP](#).

## Establishing Sessions with All VRRP Neighbors

BFD sessions can be established for all VRRP neighbors at once, or a session can be established with a particular neighbor.



**Figure 14. Establishing Sessions with All VRRP Neighbors**

To establish sessions with all VRRP neighbors, use the following command.

- Establish sessions with all VRRP neighbors.

INTERFACE mode

```
vrrp bfd all-neighbors
```

## Establishing VRRP Sessions on VRRP Neighbors

The master router does not care about the state of the backup router, so it does not participate in any VRRP BFD sessions.

VRRP BFD sessions on the backup router cannot change to the UP state. Configure the master router to establish an individual VRRP session the backup router.

To establish a session with a particular VRRP neighbor, use the following command.

- Establish a session with a particular VRRP neighbor.

INTERFACE mode

```
vrrp bfd neighbor ip-address
```

### Example of Viewing Sessions with VRRP Neighbors

#### Example of Viewing VRRP Session State Information

To view the established sessions, use the `show bfd neighbors` command.

The bold line shows that VRRP BFD sessions are enabled.

```
R1(conf-if-gi-4/25)#vrrp bfd all-neighbors
R1(conf-if-gi-4/25)#do show bfd neighbor

* - Active session role
Ad Dn - Admin Down
C - CLI
I - ISIS
O - OSPF
R - Static Route (RTM)
V - VRRP
```

```
LocalAddr RemoteAddr Interface State Rx-int Tx-int Mult Clients
* 2.2.5.1 2.2.5.2 Gi 4/25 Down 1000 1000 3 V
```

To view session state information, use the `show vrrp` command.

The bold line shows the VRRP BFD session.

```
R1(conf-if-gi-4/25)#do show vrrp

GigabitEthernet 4/1, VRID: 1, Net: 2.2.5.1
State: Backup, Priority: 1, Master: 2.2.5.2
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 95, Bad pkts rcvd: 0, Adv sent: 933, Gratuitous ARP sent: 3
Virtual MAC address:
 00:00:5e:00:01:01
Virtual IP address:
 2.2.5.4
Authentication: (none)
BFD Neighbors:
RemoteAddr State
2.2.5.2 Up
```

## Changing VRRP Session Parameters

BFD sessions are configured with default intervals and a default role.

The parameters that you can configure are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. You can change parameters for all VRRP sessions or for a particular neighbor.

To change parameters for all VRRP sessions or for a particular VRRP session, use the following commands.

- Change parameters for all VRRP sessions.

INTERFACE mode

```
vrp bfd all-neighbors interval milliseconds min_rx milliseconds multiplier
value role [active | passive]
```

- Change parameters for a particular VRRP session.

INTERFACE mode

```
vrp bfd neighbor ip-address interval milliseconds min_rx milliseconds
multiplier value role [active | passive]
```

To view session parameters, use the `show bfd neighbors detail` command, as shown in the example in [Verifying BFD Sessions with BGP Neighbors Using the show bfd neighbors command](#) example in [Displaying BFD for BGP Information](#).

## Disabling BFD for VRRP

If you disable any or all VRRP sessions, the sessions are torn down.

A final Admin Down control packet is sent to all neighbors and sessions on the remote system change to the Down state.

To disable all VRRP sessions on an interface, sessions for a particular VRRP group, or for a particular VRRP session on an interface, use the following commands.

- Disable all VRRP sessions on an interface.

INTERFACE mode

```
no vrrp bfd all-neighbors
```

- Disable all VRRP sessions in a VRRP group.

VRRP mode

```
bfd disable
```

- Disable a particular VRRP session on an interface.

INTERFACE mode

```
no vrrp bfd neighbor ip-address
```

## Configure BFD for VLANs

BFD on Dell Networking systems is a Layer 3 protocol.

Use BFD with routed virtual local area networks (VLANs). BFD on VLANs is analogous to BFD on physical ports. If you enable the no routing protocol, and a remote system fails, the local system does not remove the connected route until the first failed attempt to send a packet. If you enable BFD, the local system removes the route when it stops receiving periodic control packets from the remote system.

There is one BFD agent for VLANs and port-channels that resides on RP2, as opposed to the other agents that are on the line card. Therefore, the 100 total possible sessions that this agent can maintain is shared for VLANs and port-channels.

Configuring BFD for VLANs is a two-step process:

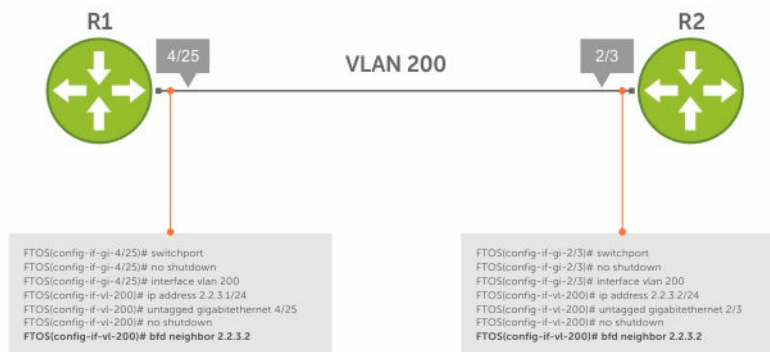
1. Enable the BFD globally. Refer to [Enabling BFD Globally](#).
2. Establish sessions with VLAN neighbors. Refer to [Establish Sessions with VLAN Neighbors](#).

## Related Configuration Task

- [Changing VLAN Session Parameters](#).
- [Disabling BFD for VLANs](#).

## Establish Sessions with VLAN Neighbors

To establish a session, enable BFD at interface level on both ends of the link, as shown in the following illustration. The session parameters do not need to match.



**Figure 15. Establishing Sessions with VLAN Neighbors**

To establish a BFD session with a VLAN neighbor, follow this step.

- Establish sessions with a VLAN neighbor.

INTERFACE VLAN mode

```
bfd neighbor ip-address
```

View the established sessions using the `show bfd neighbors` command, as shown in the following example.

```
R2(conf-if-vl-200)#bfd neighbor 2.2.3.2
R2(conf-if-vl-200)#do show bfd neighbors
```

```
* - Active session role
Ad Dn - Admin Down
C - CLI
I - ISIS
O - OSPF
R - Static Route (RTM)
V - VRRP
```

```
LocalAddr RemoteAddr Interface State Rx-int Tx-int Mult Clients
* 2.2.3.2 2.2.3.1 v1 200 Up 100 100 3 C
```

## Changing VLAN Session Parameters

BFD sessions are configured with default intervals and a default role. The parameters that you can configure are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. You can change parameters per interface, if you make a configuration change, the change affects all sessions on that interface.

**△ CAUTION: When configuring BFD on VLAN or LAG interfaces, Dell Networking recommends a minimum value of 500 milliseconds for both the transmit and minimum receive time, which yields a final detection time of (500ms \*3) 1500 milliseconds.**

To change parameters for a session, use the following commands.

- Change session parameters for all sessions on an interface.

```
INTERFACE VLAN mode
```

```
bfd interval milliseconds min_rx milliseconds multiplier value role [active | passive]
```

To view session parameters, use the `show bfd neighbors` command, as shown in the example [Changing Physical Port Session Parameters](#).

## Disabling BFD for VLANs

If you disable BFD on an interface, sessions on the interface are torn down. A final Admin Down control packet is sent to all neighbors and sessions on the remote system change to the Down state.

To disable BFD on a VLAN interface, use the following command.

- Disable all sessions on a VLAN interface.

```
INTERFACE VLAN mode
```

```
no bfd enable
```

## Configure BFD for Port-Channels

BFD on port-channels is analogous to BFD on physical ports.

If you enable the no routing protocol, and a remote system fails, the local system does not remove the connected route until the first failed attempt to send a packet. If you enable BFD, the local system removes the route when it stops receiving periodic control packets from the remote system.

There is one BFD agent for VLANs and port-channels that resides on RP2, as opposed to the other agents that are on the line card. Therefore, the 100 total possible sessions that this agent can maintain is shared for VLANs and port-channels.

Configuring BFD for port-channels is a two-step process:

- Enable BFD globally. Refer to [Enabling BFD Globally](#).

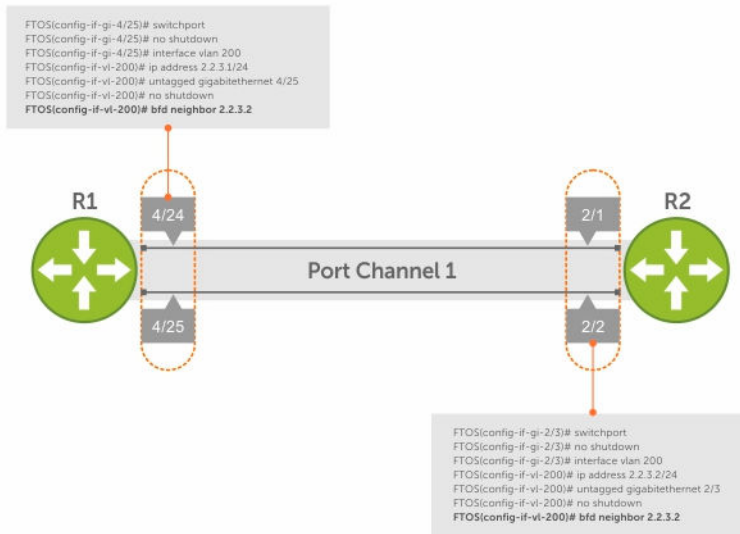
- Establish sessions on port-channels. Refer to [Establish Sessions on Port-Channels](#).

## Related Configuration Tasks

- [Changing Port-Channel Session Parameters](#).
- [Disabling BFD for Port-Channels](#).

## Establish Sessions on Port-Channels

To establish a session, you must enable BFD at interface level on both ends of the link, as shown in the following example. The session parameters do not need to match.



**Figure 16. Establishing Sessions on Port-Channels**

To establish a session on a port-channel, use the `bfd neighbor ip-address` command in INTERFACE PORT-CHANNEL mode.

View the established sessions using the `show bfd neighbors` command, as shown in [Changing Port-Channel Session Parameters](#).

### Viewing Established Sessions for VLAN Neighbors

```

R2(conf-if-po-1)#bfd neighbors 2.2.2.1
R2(conf-if-po-1)#do show bfd neighbors
* - Active session role
Ad Dn - Admin Down
C - CLI
I - ISIS
O - OSPF
R - Static Route (RTM)
V - VRRP
 LocalAddr RemoteAddr Interface State Rx-int Tx-int Mult Clients
* 2.2.2.2 2.2.2.1 Po 1 Up 100 100 3 C

```



## Changing Physical Port Session Parameters

Configure BFD sessions with default intervals and a default role.

The parameters that you can configure are: Desired TX Interval, Required Min RX Interval, Detection Multiplier, and system role. Configure these parameters per interface; if you change a parameter, the change affects all physical port sessions on that interface.

**△ CAUTION: When configuring BFD on VLAN or LAG interfaces, Dell Networking recommends a minimum value of 500 milliseconds for both the transmit and minimum receive time, which yields a final detection time of (500ms \*3) 1500 milliseconds.**

Change session parameters for all sessions on an interface.

INTERFACE PORT-CHANNEL mode

```
bfd interval milliseconds min_rx milliseconds multiplier value role [active
| passive]
```

View session parameters using the `show bfd neighbors detail` command.

## Disabling BFD for Port-Channels

If you disable BFD on an interface, sessions on the interface are torn down.

A final Admin Down control packet is sent to all neighbors, and sessions on the remote system are placed in a Down state.

To disable BFD for a port-channel, use the following command.

- Disable BFD for a port-channel.

INTERFACE PORT-CHANNEL mode

```
no bfd enable
```

## Configuring Protocol Liveness

Protocol liveness is a feature that notifies the BFD manager when a client protocol is disabled.

When you disable a client, all BFD sessions for that protocol are torn down. Neighbors on the remote system receive an Admin Down control packet and are placed in the Down state.

To enable protocol liveness, use the following command.

- Enable Protocol Liveness.

CONFIGURATION mode

```
bfd protocol-liveness
```

# Troubleshooting BFD

To troubleshoot BFD, use the following commands and examples.

To control packet field values or to examine the control packets in hexadecimal format, use the following command.

- Examine control packet field values.

```
CONFIGURATION mode
```

```
debug bfd detail
```

- Examine the control packets in hexadecimal format.

```
CONFIGURATION
```

```
debug bfd packet
```

## Example of Output from the debug bfd detail Command

## Example of Output from the debug bfd packet Command

The following example shows a three-way handshake using the debug bfd detail command.

```
R1(conf-if-gi-4/24)#00:54:38: %RPM0-P:RP2 %BFDMGR-1-BFD_STATE_CHANGE:
Changed session state to
Down for neighbor 2.2.2.2 on interface Gi 4/24 (diag: 0)
00:54:38 : Sent packet for session with neighbor 2.2.2.2 on Gi 4/24
TX packet dump:
Version:1, Diag code:0, State:Down, Poll bit:0, Final bit:0, Demand bit:0
myDiscrim:4, yourDiscrim:0, minTx:1000000, minRx:1000000, multiplier:3,
minEchoRx:0
00:54:38 : Received packet for session with neighbor 2.2.2.2 on Gi 4/24
RX packet dump:
Version:1, Diag code:0, State:Init, Poll bit:0, Final bit:0, Demand bit:0
myDiscrim:6, yourDiscrim:4, minTx:1000000, minRx:1000000, multiplier:3,
minEchoRx:0
00:54:38: %RPM0-P:RP2 %BFDMGR-1-BFD_STATE_CHANGE: Changed session state to
Up for neighbor 2.2.2.2
on interface Gi 4/24 (diag: 0)
```

The following example displays hexadecimal output from the debug bfd packet command.

```
RX packet dump:
20 c0 03 18 00 00 00 05 00 00 00 04 00 01 86 a0
00 01 86 a0 00 00 00 00
00:34:13 : Sent packet for session with neighbor 2.2.2.2 on Gi 4/24
TX packet dump:
20 c0 03 18 00 00 00 04 00 00 00 05 00 01 86 a0
00 01 86 a0 00 00 00 00
00:34:14 : Received packet for session with neighbor 2.2.2.2 on Gi 4/24
RX packet dump:
20 c0 03 18 00 00 00 05 00 00 00 04 00 01 86 a0
00 01 86 a0 00 00 00 00
00:34:14 : Sent packet for session with neighbor 2.2.2.2 on Gi 4/24
TX packet dump:
```

The output for the debug bfd event command is the same as the log messages that appear on the console by default.

# Border Gateway Protocol IPv4 (BGPv4)

This chapter provides a general description of BGPv4 as it is supported in the Dell Networking operating system.

BGP protocol standards are listed in the [Standards Compliance](#) chapter.

BGP is an external gateway protocol that transmits interdomain routing information within and between autonomous systems (AS). The primary function of the BGP is to exchange network reachability information with other BGP systems. BGP generally operates with an internal gateway protocol (IGP) such as open shortest path first (OSPF) or router information protocol (RIP), allowing you to communicate to external ASs smoothly. BGP adds reliability to network connections by having multiple paths from one router to another.

## Autonomous Systems (AS)

BGP autonomous systems (ASs) are a collection of nodes under common administration with common network routing policies.


Each AS has a number, which an internet authority already assigns. You do not assign the BGP number.

AS numbers (ASNs) are important because the ASN uniquely identifies each network on the Internet. The Internet Assigned Numbers Authority (IANA) has reserved AS numbers 64512 through 65534 to be used for private purposes. IANA reserves ASNs 0 and 65535 and must not be used in a live environment.

You can group autonomous systems into three categories (multihomed, stub, and transit), defined by their connections and operation.

- **multihomed AS** — is one that maintains connections to more than one other AS. This group allows the AS to remain connected to the Internet in the event of a complete failure of one of their connections. However, this type of AS does not allow traffic from one AS to pass through on its way to another AS. A simple example of this group is seen in the following illustration.
- **stub AS** — is one that is connected to only one other AS.
- **transit AS** — is one that provides connections through itself to separate networks. For example, in the following illustration, Router 1 can use Router 2 (the transit AS) to connect to Router 4. Internet service providers (ISPs) are always transit ASs, because they provide connections from one network to another. The ISP is considered to be “selling transit service” to the customer network, so thus the term Transit AS.

When BGP operates inside an AS (AS1 or AS2, as seen in the following illustration), it is referred to as Internal BGP (IBGP Interior Border Gateway Protocol). When BGP operates between ASs (AS1 and AS2), it is called External BGP (EBGP Exterior Border Gateway Protocol). IBGP provides routers inside the AS with the knowledge to reach routers external to the AS. EBGP routers exchange information with other EBGP routers as well as IBGP routers to maintain connectivity and accessibility.




**Figure 17. Interior BGP**

BGP version 4 (BGPv4) supports classless interdomain routing and aggregate routes and AS paths. BGP is a path vector protocol — a computer network in which BGP maintains the path that updated information takes as it diffuses through the network. Updates traveling through the network and returning to the same node are easily detected and discarded.

BGP does not use a traditional interior gateway protocol (IGP) matrix, but makes routing decisions based on path, network policies, and/or rulesets. Unlike most protocols, BGP uses TCP as its transport protocol.

Because each BGP router talking to another router is a session, a BGP network needs to be in “full mesh.” This is a topology that has every router directly connected to every other router. Each BGP router within an AS must have iBGP sessions with all other BGP routers in the AS. For example, a BGP network within an AS needs to be in “full mesh.” As seen in the following illustration, four routers connected in a full mesh have three peers each, six routers have five peers each, and eight routers in full mesh have seven peers each.



**Figure 18. BGP Routers in Full Mesh**

The number of BGP speakers each BGP peer must maintain increases exponentially. Network management quickly becomes impossible.

## Sessions and Peers

When two routers communicate using the BGP protocol, a BGP session is started. The two end-points of that session are Peers. A Peer is also called a Neighbor.

## Establish a Session

Information exchange between peers is driven by events and timers. The focus in BGP is on the traffic routing policies.

In order to make decisions in its operations with other BGP peers, a BGP process uses a simple finite state machine that consists of six states: Idle, Connect, Active, OpenSent, OpenConfirm, and Established. For each peer-to-peer session, a BGP implementation tracks which of these six states the session is in. The BGP protocol defines the messages that each peer should exchange in order to change the session from one state to another.

| State                            | Description                                                                                                                                                                                                                                                  |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Idle</b>                      | BGP initializes all resources, refuses all inbound BGP connection attempts, and initiates a TCP connection to the peer.                                                                                                                                      |
| <b>Connect</b>                   | In this state the router waits for the TCP connection to complete, transitioning to the OpenSent state if successful.<br>If that transition is not successful, BGP resets the ConnectRetry timer and transitions to the Active state when the timer expires. |
| <b>Active</b>                    | The router resets the ConnectRetry timer to zero and returns to the Connect state.                                                                                                                                                                           |
| <b>OpenSent</b>                  | After successful OpenSent transition, the router sends an Open message and waits for one in return.                                                                                                                                                          |
| <b>OpenConfirm</b>               | After the Open message parameters are agreed between peers, the neighbor relation is established and is in the OpenConfirm state. This is when the router receives and checks for agreement on the parameters of open messages to establish a session.       |
| <b>Keepalive and Established</b> | Keepalive messages are exchanged next, and after successful receipt, the router is placed in the Established state. Keepalive messages continue to be sent at regular periods (established by the Keepalive timer) to verify connections.                    |

After the connection is established, the router can now send/receive Keepalive, Update, and Notification messages to/from its peer.

## Peer Groups

Peer Groups are neighbors grouped according to common routing policies. They enable easier system configuration and management by allowing groups of routers to share and inherit policies.

Peer groups also aid in convergence speed. When a BGP process needs to send the same information to a large number of peers, the BGP process needs to set up a long output queue to get that information to all the proper peers. If the peers are members of a peer group however, the information can be sent to one place and then passed onto the peers within the group.

## Route Reflectors


Route reflectors (RR) reorganize the iBGP core into a hierarchy and allow some route advertisement rules.

Route reflection divides iBGP peers into two groups: client peers and nonclient peers. A route reflector and its client peers form a route reflection cluster. Because BGP speakers announce only the best route for a given prefix, route reflector rules are applied after the router makes its best path decision.

 **NOTE:** Address-family specific RR configurations are not supported.

- If a route was received from a nonclient peer, reflect the route to all client peers.
- If the route was received from a client peer, reflect the route to all nonclient and all client peers.

To illustrate how these rules affect routing, refer to the following illustration and the following steps. Routers B, C, D, E, and G are members of the same AS (AS100). These routers are also in the same Route Reflection Cluster, where Router D is the Route Reflector. Router E and H are client peers of Router D; Routers B and C and nonclient peers of Router D.



**Figure 19. BGP Router Rules**

1. Router B receives an advertisement from Router A through eBGP. Because the route is learned through eBGP, Router B advertises it to all its iBGP peers: Routers C and D.
2. Router C receives the advertisement but does not advertise it to any peer because its only other peer is Router D, an iBGP peer, and Router D has already learned it through iBGP from Router B.
3. Router D does not advertise the route to Router C because Router C is a nonclient peer and the route advertisement came from Router B who is also a nonclient peer.
4. Router D does reflect the advertisement to Routers E and G because they are client peers of Router D.
5. Routers E and G then advertise this iBGP learned route to their eBGP peers Routers F and H.

## Communities

BGP communities are sets of routes with one or more common attributes. Communities are a way to assign common attributes to multiple routes at the same time.

## BGP Attributes

Routes learned using BGP have associated properties that are used to determine the best route to a destination when multiple paths exist to a particular destination.

These properties are referred to as BGP attributes, and an understanding of how BGP attributes influence route selection is required for the design of robust networks. This section describes the attributes that BGP uses in the route selection process:


- [Weight](#)
- [Local Preference](#)
- [Multi-Exit Discriminators \(MEDs\)](#)
- [Origin](#)
- [AS Path](#)
- [Next Hop](#)

### Best Path Selection Criteria

Paths for active routes are grouped in ascending order according to their neighboring external AS number (BGP best path selection is deterministic by default, which means the `bgp non-deterministic-med` command is NOT applied).


The best path in each group is selected based on specific criteria. Only one “best path” is selected at a time. If any of the criteria results in more than one path, BGP moves on to the next option in the list. For example, two paths may have the same weights, but different local preferences. BGP sees that the Weight criteria results in two potential “best paths” and moves to local preference to reduce the options. If a number of best paths are determined, this selection criteria is applied to group’s best to determine the ultimate best path.

In non-deterministic mode (the `bgp non-deterministic-med` command is applied), paths are compared in the order in which they arrive. This method can lead to the system choosing different best paths from a set of paths, depending on the order in which they were received from the neighbors because MED may or may not get compared between the adjacent paths. In deterministic mode, the system compares MED between the adjacent paths within an AS group because all paths in the AS group are from the same AS.

 **NOTE:** In the Dell Networking OS version 8.3.11.4, the `bgp bestpath as-path multipath-relax` command is disabled by default, preventing BGP from load-balancing a learned route across two or more eBGP peers. To enable load-balancing across different eBGP peers, enable the `bgp bestpath as-path multipath-relax` command. A system error results if you configure the `bgp bestpath as-path ignore` command and the `bgp bestpath as-path multipath-relax` command at the same time. Only enable one command at a time.

The following illustration shows that the decisions BGP goes through to select the best path. The list following the illustration details the path selection criteria.





**Figure 20. BGP Best Path Selection**

**Best Path Selection Details**

1. Prefer the path with the largest WEIGHT attribute.
2. Prefer the path with the largest LOCAL\_PREF attribute.
3. Prefer the path that was locally Originated via a network command, redistribute command or aggregate-address command.
  - a. Routes originated with the Originated via a network OR redistribute commands are preferred over routes originated with the aggregate-address command.
4. Prefer the path with the shortest AS\_PATH (unless the `bgp bestpath as-path ignore` command is configured, then AS\_PATH is not considered). The following criteria apply:
  - a. An AS\_SET has a path length of 1, no matter how many ASs are in the set.
  - b. A path with no AS\_PATH configured has a path length of 0.
  - c. AS\_CONFED\_SET is not included in the AS\_PATH length.
  - d. AS\_CONFED\_SEQUENCE has a path length of 1, no matter how many ASs are in the AS\_CONFED\_SEQUENCE.
5. Prefer the path with the lowest ORIGIN type (IGP is lower than EGP, and EGP is lower than INCOMPLETE).
6. Prefer the path with the lowest multi-exit discriminator (MED) attribute. The following criteria apply:
  - a. This comparison is only done if the first (neighboring) AS is the same in the two paths; the MEDs are compared only if the first AS in the AS\_SEQUENCE is the same for both paths.
  - b. If you entered the `bgp always-compare-med` command, MEDs are compared for all paths.

- c. Paths with no MED are treated as “worst” and assigned a MED of 4294967295.
- 7. Prefer external (EBGP) to internal (IBGP) paths or confederation EBGP paths.
- 8. Prefer the path with the lowest IGP metric to the BGP if next-hop is selected when `synchronization` is disabled and only an internal path remains.
- 9. The system deems the paths as equal and does not perform steps 9 through 11, if the following criteria is met:
  - a. the IBGP multipath or EBGP multipath are configured (the `maximum-path` command).
  - b. the paths being compared were received from the same AS with the same number of ASs in the AS Path but with different NextHops.
  - c. the paths were received from IBGP or EBGP neighbor respectively.
- 10. If the `bgp bestpath router-id ignore` command is enabled and:
  - a. if the Router-ID is the same for multiple paths (because the routes were received from the same route) skip this step.
  - b. if the Router-ID is NOT the same for multiple paths, prefer the path that was first received as the Best Path. The path selection algorithm returns without performing any of the checks detailed here.
- 11. Prefer the external path originated from the BGP router with the lowest router ID. If both paths are external, prefer the oldest path (first received path). For paths containing a route reflector (RR) attribute, the originator ID is substituted for the router ID.
- 12. If two paths have the same router ID, prefer the path with the lowest cluster ID length. Paths without a cluster ID length are set to a 0 cluster ID length.
- 13. Prefer the path originated from the neighbor with the lowest address. (The neighbor address is used in the BGP neighbor configuration and corresponds to the remote peer used in the TCP connection with the local router.)

After a number of best paths are determined, this selection criteria is applied to group’s best to determine the ultimate best path.

In non-deterministic mode (the `bgp non-deterministic-med` command is applied), paths are compared in the order in which they arrive. This method can lead to the system choosing different best paths from a set of paths, depending on the order in which they were received from the neighbors because MED may or may not get compared between the adjacent paths. In deterministic mode, the system compares MED between the adjacent paths within an AS group because all paths in the AS group are from the same AS.

## Weight

The weight attribute is local to the router and is not advertised to neighboring routers.

If the router learns about more than one route to the same destination, the route with the highest weight is preferred. The route with the highest weight is installed in the IP routing table.

## Local Preference

Local preference (LOCAL\_PREF) represents the degree of preference within the entire AS. The higher the number, the greater the preference for the route.

Local preference (LOCAL\_PREF) is one of the criteria used to determine the best path, so keep in mind that other criteria may impact selection, as shown in the illustration in [Best Path Selection Criteria](#). For this example, assume that the local preference (LOCAL\_PREF) is the only attribute applied. In the following illustration, AS100 has two possible paths to AS 200. Although the path through Router A is

shorter (one hop instead of two), the LOCAL\_PREF settings have the preferred path go through Router B and AS300. This is advertised to all routers within AS100, causing all BGP speakers to prefer the path through Router B.




Figure 21. BGP Local Preference


## Multi-Exit Discriminators (MEDs)


If two ASs connect in more than one place, a multi-exit discriminator (MED) can be used to assign a preference to a preferred path.

MED is one of the criteria used to determine the best path, so keep in mind that other criteria may impact selection, as shown in the illustration in [Best Path Selection Criteria](#).

One AS assigns the MED a value and the other AS uses that value to decide the preferred path. For this example, assume the MED is the only attribute applied. In the following illustration, AS100 and AS200 connect in two places. Each connection is a BGP session. AS200 sets the MED for its T1 exit point to 100 and the MED for its OC3 exit point to 50. This sets up a path preference through the OC3 link. The MEDs are advertised to AS100 routers so they know which is the preferred path.

MEDs are non-transitive attributes. If AS100 sends an MED to AS200, AS200 does not pass it on to AS300 or AS400. The MED is a locally relevant attribute to the two participating ASs (AS100 and AS200).

 **NOTE:** The MEDs are advertised across both links, so if a link goes down, AS 1 still has connectivity to AS300 and AS400.



**Figure 22. Multi-Exit Discriminators**

**NOTE:** With the Dell Networking OS version 8.3.1.0, configuring the `set metric-type internal` command in a route-map advertises the IGP cost as MED to outbound EBGP peers when redistributing routes. The configured `set metric` value overwrites the default IGP cost.

## Origin

The origin indicates the origin of the prefix, or how the prefix came into BGP. There are three origin codes: IGP, EGP, INCOMPLETE.

| Origin Type | Description                                                                                        |
|-------------|----------------------------------------------------------------------------------------------------|
| IGP         | Indicates the prefix originated from information learned through an interior gateway protocol.     |
| EGP         | Indicates the prefix originated from information learned from an EGP protocol, which NGP replaced. |
| INCOMPLETE  | Indicates that the prefix originated from an unknown source.                                       |

Generally, an IGP indicator means that the route was derived inside the originating AS. EGP generally means that a route was learned from an external gateway protocol. An INCOMPLETE origin code generally results from aggregation, redistribution, or other indirect ways of installing routes into BGP.

In the Dell Networking OS, these origin codes appear as shown in the following example. The question mark (?) indicates an origin code of INCOMPLETE (shown in bold). The lower case letter (i) indicates an origin code of IGP (shown in bold).

### Example of Viewing Origin Codes

```
Dell#show ip bgp
BGP table version is 0, local router ID is 10.101.15.13
Status codes: s suppressed, d damped, h history, * valid, > best
Path source: I - internal, a - aggregate, c - confed-external, r -
redistributed, n - network
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
Network Next Hop Metric LocPrf Weight Path
```

```
*> 7.0.0.0/29 10.114.8.33 0 0 18508 ?
*> 7.0.0.0/30 10.114.8.33 0 0 18508 ?
*> 9.2.0.0/16 10.114.8.33 10 0 18508 701 i
```

## AS Path

The AS path is the list of all ASs that all the prefixes listed in the update have passed through.

The local AS number is added by the BGP speaker when advertising to a eBGP neighbor.

The AS path is shown in the following example. The origin attribute is shown following the AS path information (shown in bold).

### Example of Viewing AS Paths

```
Dell#show ip bgp paths
Total 30655 Paths
Address Hash Refcount Metric Path
0x4014154 0 3 18508 701 3549 19421 i
0x4013914 0 3 18508 701 7018 14990 i
0x5166d6c 0 3 18508 209 4637 1221 9249 9249 i
0x5e62df4 0 2 18508 701 17302 i
0x3a1814c 0 26 18508 209 22291 i
0x567ea9c 0 75 18508 209 3356 2529 i
0x6cc1294 0 2 18508 209 1239 19265 i
0x6cc18d4 0 1 18508 701 2914 4713 17935 i
0x5982e44 0 162 18508 209 i
0x67d4a14 0 2 18508 701 19878 ?
0x559972c 0 31 18508 209 18756 i
0x59cd3b4 0 2 18508 209 7018 15227 i
0x7128114 0 10 18508 209 3356 13845 i
0x536a914 0 3 18508 209 701 6347 7781 i
0x2ffe884 0 1 18508 701 3561 9116 21350 i
```

## Next Hop

The next hop is the IP address used to reach the advertising router.


For EBGP neighbors, the next-hop address is the IP address of the connection between the neighbors. For IBGP, the EBGP next-hop address is carried into the local AS. A next hop attribute is set when a BGP speaker advertises itself to another BGP speaker outside its local AS and when advertising routes within an AS. The next hop attribute also serves as a way to direct traffic to another BGP speaker, rather than waiting for a speaker to advertise.

The system allows you to set the next hop attribute in the CLI. Setting the next hop attribute lets you determine a router as the next hop for a BGP neighbor.

## Multiprotocol BGP

Multiprotocol extensions for BGP (MBGP) is defined in IETF RFC 2858. MBGP allows different types of address families to be distributed in parallel.

MBGP allows information about the topology of the IP multicast-capable routers to be exchanged separately from the topology of normal IPv4 and IPv6 unicast routers. It allows a multicast routing topology different from the unicast routing topology.

 **NOTE:** It is possible to configure BGP peers that exchange both unicast and multicast network layer reachability information (NLR), but you cannot connect multiprotocol BGP with BGP. Therefore, you cannot redistribute multiprotocol BGP routes into BGP.

## Implement BGP with the Dell Networking OS

The following sections describe how to implement BGP on the Dell Networking OS.

### Additional Path (Add-Path) Support

The add-path feature reduces convergence times by advertising multiple paths to its peers for the same address prefix without replacing existing paths with new ones. By default, a BGP speaker advertises only the best path to its peers for a given address prefix. If the best path becomes unavailable, the BGP speaker withdraws its path from its local RIB and recalculates a new best path. This situation requires both IGP and BGP convergence and can be a lengthy process.

BGP add-path reduces the time taken for BGP convergence by advertising multiple paths to its peers for the same address prefix without new paths implicitly replacing the existing paths. An iBGP speaker that receives multiple paths from its peers should calculate the best path in its own. BGP add-path helps switchover to next new best path based on IGP convergence time when best path becomes unavailable.

### Advertise IGP Cost as MED for Redistributed Routes

When using multipath connectivity to an external AS, you can advertise the MED value selectively to each peer for redistributed routes. For some peers you can set the internal/IGP cost as the MED while setting others to a constant pre-defined metric as MED value.

The Dell Networking OS version 8.3.1.0 and later support configuring the `set metric-type internal` command in a route-map to advertise the IGP cost as the MED to outbound EBGP peers when redistributing routes. The configured `set metric` value overwrites the default IGP cost.


By using the `redistribute` command with the `route-map` command, you can specify whether a peer advertises the standard MED or uses the IGP cost as the MED.

When configuring this functionality:

- If the `redistribute` command does not have `metric` configured and the BGP peer outbound route-map does have `metric-type internal` configured, BGP advertises the IGP cost as MED.
- If the `redistribute` command has `metric` configured (`route-map set metric or redistribute route-type metric`) and the BGP peer outbound route-map has `metric-type`

internal configured, BGP advertises the metric configured in the `redistribute` command as MED.

- If BGP peer outbound route-map has `metric` configured, all other metrics are overwritten by this configuration.

 **NOTE:** When redistributing static, connected, or OSPF routes, there is no `metric` option. Simply assign the appropriate route-map to the redistributed route.

The following table lists some examples of these rules.

**Table 6. Redistributed Route Rules**

| Command Settings                                       | BGP Local Routing Information Base | MED Advertised to Peer WITH route-map metric-type internal | MED Advertised to Peer WITHOUT route-map metric-type internal |
|--------------------------------------------------------|------------------------------------|------------------------------------------------------------|---------------------------------------------------------------|
| <code>redistribute isis (IGP cost = 20)</code>         | MED: IGP cost 20                   | MED = 20                                                   | MED = 0                                                       |
| <code>redistribute isis route-map set metric 50</code> | MED: IGP cost 50                   | MED: 50 MED: 50                                            | MED: 50 MED: 50                                               |
| <code>redistribute isis metric 100</code>              | MED: IGP cost 100                  | MED: 100                                                   | MED: 100                                                      |

## Ignore Router-ID for Some Best-Path Calculations

The Dell Networking OS version 8.3.1.0 and later allows you to avoid unnecessary BGP best-path transitions between external paths under certain conditions. The `bgp bestpath router-id ignore` command reduces network disruption caused by routing and forwarding plane changes and allows for faster convergence.

## Four-Byte AS Numbers

The Dell Networking OS version 7.7.1 and later supports 4-Byte (32-bit) format when configuring autonomous system numbers (ASNs).

The 4-Byte support is advertised as a new BGP capability (4-BYTE-AS) in the OPEN message. If a 4-Byte BGP speaker has sent and received this capability from another speaker, all the messages will be 4-octet. The behavior of a 4-Byte BGP speaker is different with the peer depending on whether the peer is a 4-Byte or 2-Byte BGP speaker.

Where the 2-Byte format is 1-65535, the 4-Byte format is 1-4294967295. Enter AS numbers using the traditional format. If the ASN is greater than 65535, the dot format is shown when using the `show ip bgp` commands. For example, an ASN entered as 3183856184 appears in the `show` commands as 48581.51768; an ASN of 65123 is shown as 65123. To calculate the comparable dot format for an ASN from a traditional format, use `ASN/65536`. `ASN%65536`.

| Traditional Format | DOT Format |
|--------------------|------------|
| 65001              | 0.65501    |
| 65536              | 1.0        |
| 100000             | 1.34464    |

**Traditional Format**    **DOT Format**


**4294967295**            65535.65535

When creating Confederations, all the routers in a Confederation must be either 4-Byte or 2-Byte identified routers. You cannot mix them.

Configure 4-byte AS numbers with the `four-octet-support` command.

## AS4 Number Representation

The Dell Networking OS version 8.2.1.0 supports multiple representations of 4-byte AS numbers: `asplain`, `asdot+`, and `asdot`.

 **NOTE:** The `ASDOT` and `ASDOT+` representations are supported only with the 4-Byte AS numbers feature. If 4-Byte AS numbers are not implemented, only `ASPLAIN` representation is supported.

`ASPLAIN` is the method the Dell Networking OS has used for all previous Dell Networking OS versions. `ASPLAIN` remains the default method with the Dell Networking OS version 8.2.1.0 and later. With the `ASPLAIN` notation, a 32-bit binary AS number is translated into a decimal value.

- All AS numbers between 0 and 65535 are represented as a decimal number when entered in the CLI and when displayed in the `show` commands output.
- AS numbers larger than 65535 are represented using `ASPLAIN` notation. When entered in the CLI and when displayed in the `show` commands output, 65546 is represented as 65546.

`ASDOT+` representation splits the full binary 4-byte AS number into two words of 16 bits separated by a decimal point (.): `<high-order 16 bit value>.<low-order 16 bit value>`. Some examples are shown in the following table.

- All AS numbers between 0 and 65535 are represented as a decimal number, when entered in the CLI and when displayed in the `show` commands outputs.
- AS Numbers larger than 65535 is represented using `ASDOT` notation as `<higher 2 bytes in decimal>.<lower 2 bytes in decimal>`. For example: AS 65546 is represented as 1.10.

`ASDOT` representation combines the `ASPLAIN` and `ASDOT+` representations. AS numbers less than 65536 appear in integer format (`asplain`); AS numbers equal to or greater than 65536 appear in the decimal format (`asdot+`). For example, the AS number 65526 appears as 65526 and the AS number 65546 appears as 1.10.

## Dynamic AS Number Notation Application

The Dell Networking OS version 8.3.1.0 applies the ASN notation type change dynamically to the running-config statements.

When you apply or change an `asnotation`, the type selected is reflected immediately in the running-configuration and the `show` commands (refer to the following two examples).

### Example of Dynamic Changes in the Running Configuration When Using the `bgp asnotation` Command

#### Example of the Running Configuration When AS Notation is Disabled

```
ASDOT
Dell(conf-router_bgp)#bgp asnotation asdot
Dell(conf-router_bgp)#show conf
```



```

!
router bgp 100
bgp asnotation asdot
bgp four-octet-as-support
neighbor 172.30.1.250 local-as 65057
<output truncated>

Dell(conf-router_bgp)#do show ip bgp
BGP table version is 24901, local router ID is 172.30.1.57
<output truncated>

ASDOT+
Dell(conf-router_bgp)#bgp asnotation asdot+
Dell(conf-router_bgp)#show conf
!
router bgp 100
 bgp asnotation asdot
 bgp four-octet-as-support
neighbor 172.30.1.250 local-as 65057
<output truncated>

Dell(conf-router_bgp)#do show ip bgp
BGP table version is 24901, local router ID is 172.30.1.57
<output truncated>

AS-PLAIN
Dell(conf-router_bgp)#bgp asnotation asplain+
Dell(conf-router_bgp)#sho conf
!
router bgp 100
 bgp four-octet-asdot+
 bgp four-octet-as-support
neighbor 172.30.1.250 local-as 65057
<output truncated>

Dell(conf-router_bgp)#do sho ip bgp
BGP table version is 31571, local router ID is 172.30.1.57
<output truncated>

AS-PLAIN
Dell(conf-router_bgp)#bgp asnotation asplain
Dell(conf-router_bgp)#sho conf

AS NOTATION DISABLED
Dell(conf-router_bgp)#no bgp asnotation
Dell(conf-router_bgp)#sho conf
!
router bgp 100
bgp four-octet-as-support
neighbor 172.30.1.250 local-as 65057
<output truncated>

Dell(conf-router_bgp)#do sho ip bgp
BGP table version is 28093, local router ID is 172.30.1.57

AS4 SUPPORT DISABLED
Dell(conf-router_bgp)#no bgp four-octet-as-support
Dell(conf-router_bgp)#sho conf
!
router bgp 100

```

## AS Number Migration

With this feature you can transparently change the AS number of an entire BGP network and ensure that the routes are propagated throughout the network while the migration is in progress.

When migrating one AS to another, perhaps combining ASs, an eBGP network may lose its routing to an iBGP if the ASN changes. Migration can be difficult as all the iBGP and eBGP peers of the migrating network must be updated to maintain network reachability. Essentially, Local-AS provides a capability to the BGP speaker to operate as if it belongs to "virtual" AS network besides its physical AS network.

The following illustration shows a scenario where Router A, Router B, and Router C belong to AS 100, 200, and 300, respectively. Router A acquired Router B; Router B has Router C as its customer. When Router B is migrating to Router A, it must maintain the connection with Router C without immediately updating Router C's configuration. Local-AS allows this behavior to happen by allowing Router B to appear as if it still belongs to Router B's old network (AS 200) as far as communicating with Router C is concerned.




Figure 23. Before and After AS Number Migration with Local-AS Enabled

When you complete your migration, and you have reconfigured your network with the new information, disable this feature.


If you use the `no prepend` option, the Local-AS does not prepend to the updates received from the eBGP peer. If you do not select `no prepend` (the default), the Local-AS is added to the first AS segment in the AS-PATH. If an inbound route-map is used to prepend the as-path to the update from the peer, the Local-AS is added first. For example, consider the topology described in the previous illustration. If Router B has an inbound route-map applied on Router C to prepend "65001 65002" to the as-path, the following events take place on Router B:

1. Receive and validate the update.
2. Prepend local-as 200 to as-path.
3. Prepend "65001 65002" to as-path.

Local-AS is prepended before the route-map to give an impression that update passed through a router in AS 200 before it reached Router B.

## BGP4 Management Information Base (MIB)

The FORCE10-BGP4-V2-MIB enhances Dell Networking OS BGP management information base (MIB) support with many new simple network management protocol (SNMP) objects and notifications (traps) defined in *draft-ietf-idr-bgp4-mibv2-05*. To see these enhancements, download the MIB from the Dell website.

 **NOTE:** For the *Force10-BGP4-V2-MIB* and other MIB documentation, refer to the Dell iSupport web page.

## Important Points to Remember

- The *f10BgpM2AsPathTableEntry* table, *f10BgpM2AsPathSegmentIndex*, and *f10BgpM2AsPathElementIndex* are used to retrieve a particular ASN from the AS path. These indices are assigned to the AS segments and individual ASN in each segment starting from 0. For example, an AS path list of {200 300 400} 500 consists of two segments: {200 300 400} with segment index 0 and 500 with segment index 1. ASN 200, 300, and 400 are assigned 0, 1, and 2 element indices in that order.
- Unknown optional transitive attributes within a given path attribute (PA) are assigned indices in order. These indices correspond to the *f10BgpM2PathAttrUnknownIndex* field in the *f10BgpM2PathAttrUnknownEntry* table.
- Negotiation of multiple instances of the same capability is not supported. *F10BgpM2PeerCapAnnouncedIndex* and *f10BgpM2PeerCapReceivedIndex* are ignored in the peer capability lookup.
- Configure inbound BGP soft-reconfiguration on a peer for *f10BgpM2PrefixInPrefixesRejected* to display the number of prefixes filtered due to a policy. If you do enable `BGP soft-reconfig`, the denied prefixes are not accounted for.
- *F10BgpM2AdjRibsOutRoute* stores the pointer to the NLRI in the peer's Adj-Rib-Out.
- PA Index (*f10BgpM2PathAttrIndex* field in various tables) is used to retrieve specific attributes from the PA table. The Next-Hop, RR Cluster-list, and Originator ID attributes are not stored in the PA Table and cannot be retrieved using the `index passed in` command. These fields are not populated in *f10BgpM2PathAttrEntry*, *f10BgpM2PathAttrClusterEntry*, and *f10BgpM2PathAttrOriginatorIdEntry*.
- *F10BgpM2PathAttrUnknownEntry* contains the optional-transitive attribute details.

- Query for *f10BgpM2LinkLocalNextHopEntry* returns the default value for Link-local Next-hop.
- RFC 2545 and the *f10BgpM2Rfc2545Group* are not supported.
- An SNMP query displays up to 89 AS paths. A query for a larger AS path count displays as "..." at the end of the output.
- SNMP set for BGP is not supported. For all peer configuration tables (*f10BgpM2PeerConfigurationGroup*, *f10BgpM2PeerRouteReflectorCfgGroup*, and *f10BgpM2PeerAsConfederationCfgGroup*), an SNMP set operation returns an error. Only SNMP queries are supported. In addition, the *f10BgpM2CfgPeerError*, *f10BgpM2CfgPeerBgpPeerEntry*, and *f10BgpM2CfgPeerRowEntryStatus* fields are to hold the SNMP set status and are ignored in SNMP query.
- The AFI/SAFI is not used as an index to the *f10BgpM2PeerCountersEntry* table. The BGP peer's AFI/SAFI (IPv4 Unicast or IPv6 Multicast) is used for various outbound counters. Counters corresponding to IPv4 Multicast cannot be queried.
- The *f10BgpM2[Cfg]PeerReflectorClient* field is populated based on the assumption that route-reflector clients are not in a full mesh if you enable BGP `client-2-client` reflection and that the BGP speaker acting as reflector advertises routes learned from one client to another client. If disabled, it is assumed that clients are in a full mesh and there is no need to advertise prefixes to the other clients.
- High CPU utilization may be observed during an SNMP walk of a large BGP Loc-RIB.
- To avoid SNMP timeouts with a large-scale configuration (large number of BGP neighbors and a large BGP Loc-RIB), Dell Networking recommends setting the timeout and retry count values to a relatively higher number. For example, `t = 60` or `r = 5`.
- To return all values on an snmpwalk for the *f10BgpM2Peer sub-OID*, use the `-C c` option, such as `snmpwalk -v 2c -C c -c public<IP_address><OID>`.
- An SNMP walk may terminate pre-maturely if the index does not increment lexicographically. Dell Networking recommends using options to ignore such errors.
- Multiple BGP process instances are not supported. Thus, the *f10BgpM2PeerInstance* field in various tables is not used to locate a peer.
- Multiple instances of the same NLRI in the BGP RIB are not supported and are set to zero in the SNMP query response.
- The *f10BgpM2NlriIndex* and *f10BgpM2AdjRibsOutIndex* fields are not used.
- Carrying MPLS labels in BGP is not supported. The *f10BgpM2NlriOpaqueType* and *f10BgpM2NlriOpaquePointer* fields are set to zero.
- 4-byte ASN is supported. The *f10BgpM2AsPath4byteEntry* table contains 4-byte ASN-related parameters based on the configuration.

Traps (notifications) specified in the BGP4 MIB draft `<draft-ietf-idr-bgp4-mibv2-05.txt>` are not supported. Such traps (*bgpM2Established* and *bgpM2BackwardTransition*) are supported as part of RFC 1657.

## Configuration Information

The software supports BGPv4 as well as the following:

- deterministic multi-exit discriminator (MED) (default)
- a path with a missing MED is treated as worst path and assigned an MED value of (0xffffffff)
- the community format follows RFC 1998
- delayed configuration (the software at system boot reads the entire configuration file prior to sending messages to start BGP peer sessions)

The following are not yet supported:


- auto-summarization (the default is no auto-summary)
- synchronization (the default is no synchronization)

## BGP Configuration

To enable the BGP process and begin exchanging information, assign an AS number and use commands in ROUTER BGP mode to configure a BGP neighbor.

By default, BGP is disabled.

By default, the system compares the MED attribute on different paths from within the same AS (the `bgp always-compare-med` command is not enabled).

 **NOTE:** In the Dell Networking OS, all newly configured neighbors and peer groups are disabled. To enable a neighbor or peer group, enter the `neighbor {ip-address | peer-group-name} no shutdown` command.

The following table displays the default values for BGP.

**Table 7. BGP Default Values**

| Item                           | Default                                                                                    |
|--------------------------------|--------------------------------------------------------------------------------------------|
| BGP Neighbor Adjacency changes | All BGP neighbor changes are logged.                                                       |
| Fast External Fallover feature | Disabled                                                                                   |
| Graceful Restart feature       | Disabled                                                                                   |
| Local preference               | 100                                                                                        |
| MED                            | 0                                                                                          |
| Route Flap Damping Parameters  | half-life = 15 minutes<br>reuse = 750<br>suppress = 2000<br>max-suppress-time = 60 minutes |
| Distance                       | external distance = 20<br>internal distance = 200<br>local distance = 200                  |
| Timers                         | keepalive = 60 seconds<br>holdtime = 180 seconds                                           |
| Add-path                       | Disabled                                                                                   |


## Enabling BGP

By default, BGP is not enabled on the system. The Dell Networking OS supports one autonomous system (AS) and assigns the AS number (ASN).

To establish BGP sessions and route traffic, configure at least one BGP neighbor or peer.

In BGP, routers with an established TCP connection are called neighbors or peers. After a connection is established, the neighbors exchange full BGP routing tables with incremental updates afterward. In addition, neighbors exchange KEEPALIVE messages to maintain the connection.

In BGP, neighbor routers or peers can be classified as external. External BGP peers must be connected physically to one another (unless you enable the EBGP multihop feature), while internal BGP peers do not need to be directly connected. The IP address of an EBGP neighbor is usually the IP address of the interface directly connected to the router. First, the BGP process determines if all internal BGP peers are reachable, then it determines which peers outside the AS are reachable.

 **NOTE:** Find [Sample Configurations](#) for enabling BGP routers at the end of this chapter.


1. Assign an AS number and enter ROUTER BGP mode.

CONFIGURATION mode


```
router bgp as-number
```

- *as-number*: from 0 to 65535 (2 Byte) or from 1 to 4294967295 (4 Byte) or 0.1 to 65535.65535 (Dotted format).

Only one AS is supported per system.


 **NOTE:** If you enter a 4-Byte AS number, 4-Byte AS support is enabled automatically.

- a. Enable 4-Byte support for the BGP process.

 **NOTE:** This command is OPTIONAL. Enable if you want to use 4-Byte AS numbers or if you support AS4 number representation.

CONFIG-ROUTER-BGP mode

```
bgp four-octet-as-support
```

 **NOTE:** Use it only if you support 4-Byte AS numbers or if you support AS4 number representation. If you are supporting 4-Byte ASNs, enable this command.

Disable 4-Byte support and return to the default 2-Byte format by using the `no bgp four-octet-as-support` command. You cannot disable 4-Byte support if you currently have a 4-Byte ASN configured.

Disabling 4-Byte AS numbers also disables ASDOT and ASDOT+ number representation. All AS numbers are displayed in ASPLAIN format.

- b. Enable IPv4 multicast or IPv6 mode.

CONFIG-ROUTER-BGP mode

```
address-family [ipv4 | ipv6]
```

Use this command to enter BGP for IPv6 mode (CONF-ROUTER\_BGPv6\_AF).

2. Add a neighbor as a remote AS.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group name} remote-as as-number
```

- *peer-group name*: 16 characters
- *as-number*: from 0 to 65535 (2 Byte) or from 1 to 4294967295 (4 Byte) or 0.1 to 65535.65535 (Dotted format)

Formats: IP Address A.B.C.D

You must use [Configuring Peer Groups](#) before assigning them a remote AS.

3. Enable the BGP neighbor.

CONFIG-ROUTER-BGP mode


```
neighbor {ip-address | peer-group-name} no shutdown
```

**Example of the show ip bgp summary Command (2-Byte AS number displayed)**

**Example of the show ip bgp summary Command (4-Byte AS number displayed)**

**Example of the show ip bgp neighbors Command**

**Example of Verifying BGP Configuration**

 **NOTE:** When you change the configuration of a BGP neighbor, always reset it by entering the `clear ip bgp` command in EXEC Privilege mode.

To view the BGP configuration, enter `show config` in CONFIGURATION ROUTER BGP mode. To view the BGP status, use the `show ip bgp summary` command in EXEC Privilege mode. The first example shows the summary with a 2-byte AS number displayed (in bold); the second example shows that the summary with a 4-byte AS number using the `show ip bgp summary` command (displays a 4-byte AS number in bold).

```
R2#show ip bgp summary
BGP router identifier 192.168.10.2, local AS number 65123
BGP table version is 1, main routing table version 1
1 network entrie(s) using 132 bytes of memory
1 paths using 72 bytes of memory
BGP-RIB over all using 73 bytes of memory
1 BGP path attribute entrie(s) using 72 bytes of memory
1 BGP AS-PATH entrie(s) using 47 bytes of memory
5 neighbor(s) using 23520 bytes of memory

Neighbor AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/Pfx
10.10.21.1 65123 0 0 0 0 0 never Active
10.10.32.3 65123 0 0 0 0 0 never Active
```

```
R2#show ip bgp summary
BGP router identifier 192.168.10.2, local AS number 48735.59224
BGP table version is 1, main routing table version 1
1 network entrie(s) using 132 bytes of memory
1 paths using 72 bytes of memory
BGP-RIB over all using 73 bytes of memory
1 BGP path attribute entrie(s) using 72 bytes of memory
```

```
1 BGP AS-PATH entrie(s) using 47 bytes of memory
5 neighbor(s) using 23520 bytes of memory
```

| Neighbor   | AS    | MsgRcvd | MsgSent | TblVer | InQ | OutQ | Up/Down | State/Pfx |
|------------|-------|---------|---------|--------|-----|------|---------|-----------|
| 10.10.21.1 | 65123 | 0       | 0       | 0      | 0   | 0    | never   | Active    |
| 10.10.32.3 | 65123 | 0       | 0       | 0      | 0   | 0    | never   | Active    |

For the router's identifier, the system uses the highest IP address of the Loopback interfaces configured. Because Loopback interfaces are virtual, they cannot go down, thus preventing changes in the router ID. If you do not configure Loopback interfaces, the highest IP address of any interface is used as the router ID.

To view the status of BGP neighbors, use the `show ip bgp neighbors` command in EXEC Privilege mode as shown in the first example. For BGP neighbor configuration information, use the `show running-config bgp` command in EXEC Privilege mode as shown in the second example.

 **NOTE:** The `showconfig` command in CONFIGURATION ROUTER BGP mode gives the same information as the `show running-config bgp` command.

The following example displays two neighbors: one is an external internal BGP neighbor and the second one is an internal BGP neighbor. The first line of the output for each neighbor displays the AS number and states whether the link is an external or internal (shown in bold).

The third line of the `show ip bgp neighbors` output contains the BGP State. If anything other than ESTABLISHED is listed, the neighbor is not exchanging information and routes. For more information about using the `show ip bgp neighbors` command, refer to the *Dell Networking OS Command Line Interface Reference Guide*.

```
Dell#show ip bgp neighbors
```

**BGP neighbor is 10.114.8.60, remote AS 18508, external link**

```
BGP version 4, remote router ID 10.20.20.20
BGP state ESTABLISHED, in this state for 00:01:58
Last read 00:00:14, hold time is 90, keepalive interval is 30 seconds
Received 18552 messages, 0 notifications, 0 in queue
Sent 11568 messages, 0 notifications, 0 in queue
Received 18549 updates, Sent 11562 updates
Minimum time between advertisement runs is 30 seconds
```

```
For address family: IPv4 Unicast
BGP table version 216613, neighbor version 201190
130195 accepted prefixes consume 520780 bytes
Prefix advertised 49304, rejected 0, withdrawn 36143
```

```
Connections established 1; dropped 0
Last reset never
```

```
Local host: 10.114.8.39, Local port: 1037
Foreign host: 10.114.8.60, Foreign port: 179
```

**BGP neighbor is 10.1.1.1, remote AS 65535, internal link**

```
Administratively shut down
BGP version 4, remote router ID 10.0.0.0
BGP state IDLE, in this state for 17:12:40
Last read 17:12:40, hold time is 180, keepalive interval is 60 seconds
Received 0 messages, 0 notifications, 0 in queue
Sent 0 messages, 0 notifications, 0 in queue
Received 0 updates, Sent 0 updates
Minimum time between advertisement runs is 5 seconds
```



```

For address family: IPv4 Unicast
BGP table version 0, neighbor version 0
0 accepted prefixes consume 0 bytes
Prefix advertised 0, rejected 0, withdrawn 0

Connections established 0; dropped 0
Last reset never
No active TCP connection
Dell#

```

```


R2#show running-config bgp
!
router bgp 65123
 bgp router-id 192.168.10.2
 network 10.10.21.0/24
 network 10.10.32.0/24
 network 100.10.92.0/24
 network 192.168.10.0/24
 bgp four-octet-as-support
 neighbor 10.10.21.1 remote-as 65123
 neighbor 10.10.21.1 filter-list ISP1in
 neighbor 10.10.21.1 no shutdown
 neighbor 10.10.32.3 remote-as 65123
 neighbor 10.10.32.3 no shutdown
 neighbor 100.10.92.9 remote-as 65192
 neighbor 100.10.92.9 no shutdown
 neighbor 192.168.10.1 remote-as 65123
 neighbor 192.168.10.1 update-source Loopback 0

```

## Configuring AS4 Number Representations

Enable one type of AS number representation: ASPLAIN, ASDOT+, or ASDOT.

| Term    | Description                                                                                                                                                                                                                                                                                                   |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ASPLAIN | the method Dell Networking OS used for all previous Dell Networking OS versions. It remains the default method with the Dell Networking OS version 8.2.1.0 and later. With the ASPLAIN notation, a 32-bit binary AS number is translated into a decimal value.                                                |
| ASDOT+  | representation splits the full binary 4-byte AS number into two words of 16 bits separated by a decimal point (.): <high-order 16 bit value>.<low-order 16 bit value>.                                                                                                                                        |
| ASDOT   | representation combines the ASPLAIN and ASDOT+ representations. AS numbers less than 65536 appear in integer format (asplain); AS numbers equal to or greater than 65536 appear using the decimal method (asdot+). For example, the AS number 65526 appears as 65526 and the AS number 65546 appears as 1.10. |

 **NOTE:** The ASDOT and ASDOT+ representations are supported only with the 4-Byte AS numbers feature. If you do not implement 4-Byte AS numbers, only ASPLAIN representation is supported.

Only one form of AS number representation is supported at a time. You cannot combine the types of representations within an AS.

To configure AS4 number representations, use the following commands.

- Enable ASPLAIN AS Number representation.  
CONFIG-ROUTER-BGP mode

```
bgp asnotation asplain
```



**NOTE:** ASPLAIN is the default method the system uses and does not appear in the configuration display.

- Enable ASDOT AS Number representation.  
CONFIG-ROUTER-BGP mode

```
bgp asnotation asdot
```

- Enable ASDOT+ AS Number representation.  
CONFIG-ROUTER-BGP mode

```
bgp asnotation asdot+
```

#### **Example of the `bgp asnotation asplain` Command**

#### **Example of the `bgp asnotation asdot` Command**

#### **Example of the `bgp asnotation asdot+` Command**

```
Dell(conf-router_bgp)#bgp asnotation asplain
Dell(conf-router_bgp)#sho conf
!
router bgp 100
 bgp four-octet-as-support
 neighbor 172.30.1.250 remote-as 18508
 neighbor 172.30.1.250 local-as 65057
 neighbor 172.30.1.250 route-map rmap1 in
 neighbor 172.30.1.250 password 7
5ab3eb9a15ed02ff4f0dfd4500d6017873cfd9a267c04957
```

```
Dell(conf-router_bgp)#bgp asnotation asdot
Dell(conf-router_bgp)#sho conf
!
router bgp 100
 bgp asnotation asdot
bgp four-octet-as-support
 neighbor 172.30.1.250 remote-as 18508
 neighbor 172.30.1.250 local-as 65057
 neighbor 172.30.1.250 route-map rmap1 in
 neighbor 172.30.1.250 password 7
5ab3eb9a15ed02ff4f0dfd4500d6017873cfd9a267c04957
```


```
Dell(conf-router_bgp)#bgp asnotation asdot+
Dell(conf-router_bgp)#sho conf
!
router bgp 100
 bgp asnotation asdot+
bgp four-octet-as-support
 neighbor 172.30.1.250 remote-as 18508
 neighbor 172.30.1.250 local-as 65057
 neighbor 172.30.1.250 route-map rmap1 in
 neighbor 172.30.1.250 password 7
5ab3eb9a15ed02ff4f0dfd4500d6017873cfd9a267c04957
```

## **Configuring Peer Groups**

To configure multiple BGP neighbors at one time, create and populate a BGP peer group. An advantage of peer groups is that members of a peer group inherit the configuration properties of the group and share same update policy.

A maximum of 256 peer groups are allowed on the system.

Create a peer group by assigning it a name, then adding members to the peer group. After you create a peer group, you can configure route policies for it. For information about configuring route policies for a peer group, refer to [Filtering BGP Routes](#).

 **NOTE:** Find [Sample Configurations](#) for enabling peer groups at the end of this chapter.

1. Create a peer group by assigning a name to it.  
CONFIG-ROUTERBGP mode

```
neighbor peer-group-name peer-group
```

2. Enable the peer group.  
CONFIG-ROUTERBGP mode

```
neighbor peer-group-name no shutdown
```

By default, all peer groups are disabled.

3. Create a BGP neighbor.  
CONFIG-ROUTERBGP mode

```
neighbor ip-address remote-as as-number
```

4. Enable the neighbor.  
CONFIG-ROUTERBGP mode

```
neighbor ip-address no shutdown
```

5. Add an enabled neighbor to the peer group.  
CONFIG-ROUTERBGP mode

```
neighbor ip-address peer-group peer-group-name
```

6. Add a neighbor as a remote AS.  
CONFIG-ROUTERBGP mode

```
neighbor {ip-address | peer-group name} remote-as as-number
```

Formats: IP Address A.B.C.D

- *Peer-Group Name*: 16 characters.
- *as-number*: the range is from 0 to 65535 (2-Byte) or 1 to 4294967295 | 0.1 to 65535.65535 (4-Byte) or 0.1 to 65535.65535 (Dotted format)

To add an external BGP (EBGP) neighbor, configure the *as-number* parameter with a number different from the BGP *as-number* configured in the `router bgp as-number` command.

To add an internal BGP (IBGP) neighbor, configure the *as-number* parameter with the same BGP *as-number* configured in the `router bgp as-number` command.

## Example of Viewing a Newly Created Peer Group

### Example of Enabling a Peer Group

#### Example of the `show ip bgp peer-group` Command


After you create a peer group, you can use any of the commands beginning with the keyword `neighbor` to configure that peer group.

When you add a peer to a peer group, it inherits all the peer group's configured parameters.

A neighbor cannot become part of a peer group if it has any of the following commands configured:

- `neighbor advertisement-interval`
- `neighbor distribute-list out`
- `neighbor filter-list out`
- `neighbor next-hop-self`
- `neighbor route-map out`
- `neighbor route-reflector-client`
- `neighbor send-community`

A neighbor may keep its configuration after it was added to a peer group if the neighbor's configuration is more specific than the peer group's and if the neighbor's configuration does not affect outgoing updates.

 **NOTE:** When you configure a new set of BGP policies for a peer group, *always* reset the peer group by entering the `clear ip bgp peer-group peer-group-name` command in EXEC Privilege mode.

To view the configuration, use the `show config` command in CONFIGURATION ROUTER BGP mode. When you create a peer group, it is disabled (shutdown). The following example shows the creation of a peer group (zanzibar) (in bold).

```
Dell(conf-router_bgp)#neighbor zanzibar peer-group
Dell(conf-router_bgp)#show conf
!
router bgp 45
 bgp fast-external-fallover
 bgp log-neighbor-changes
 neighbor zanzibar peer-group
 neighbor zanzibar shutdown
 neighbor 10.1.1.1 remote-as 65535
 neighbor 10.1.1.1 shutdown
```

To enable a peer group, use the `neighbor peer-group-name no shutdown` command in CONFIGURATION ROUTER BGP mode (shown in bold).

```
Dell(conf-router_bgp)#neighbor zanzibar no shutdown
Dell(conf-router_bgp)#show config
!
router bgp 45
 bgp fast-external-fallover
 bgp log-neighbor-changes
 neighbor zanzibar peer-group
 neighbor zanzibar no shutdown
 neighbor 10.1.1.1 remote-as 65535
 neighbor 10.1.1.1 shutdown
```

To disable a peer group, use the `neighbor peer-group-name shutdown` command in CONFIGURATION ROUTER BGP mode. The configuration of the peer group is maintained, but it is not applied to the peer group members. When you disable a peer group, all the peers within the peer group that are in the ESTABLISHED state move to the IDLE state.

To view the status of peer groups, use the `show ip bgp peer-group` command in EXEC Privilege mode, as shown in the following example.

```
Dell>show ip bgp peer-group

Peer-group zanzibar, remote AS 65535
BGP version 4
Minimum time between advertisement runs is 5 seconds

For address family: IPv4 Unicast
BGP neighbor is zanzibar, peer-group internal,
Number of peers in this group 26
Peer-group members (* - outbound optimized):
 10.68.160.1
 10.68.161.1
 10.68.162.1
 10.68.163.1
 10.68.164.1
 10.68.165.1
 10.68.166.1
 10.68.167.1
 10.68.168.1
 10.68.169.1
 10.68.170.1
 10.68.171.1
 10.68.172.1
 10.68.173.1
 10.68.174.1
 10.68.175.1
 10.68.176.1
 10.68.177.1
 10.68.178.1
```

## Configuring BGP Fast Fail-Over

By default, a BGP session is governed by the hold time.

BGP routers typically carry large routing tables, so frequent session resets are not desirable. The BGP fast fail-over feature reduces the convergence time while maintaining stability. The connection to a BGP peer is immediately reset if a link to a directly connected external peer fails.

When you enable fail-over, BGP tracks IP reachability to the peer remote address and the peer local address. Whenever either address becomes unreachable (for example, no active route exists in the routing table for peer IPv6 destinations/local address), BGP brings down the session with the peer.

The BGP fast fail-over feature is configured on a per-neighbor or peer-group basis and is disabled by default.

To enable the BGP fast fail-over feature, use the following command.

To disable fast fail-over, use the `[no] neighbor [neighbor | peer-group] fail-over` command in CONFIGURATION ROUTER BGP mode.

- Enable BGP Fast Fail-Over.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} fail-over
```

### Example of Verifying that Fast Fail-Over is Enabled on a BGP Neighbor

#### Example of Verifying that Fast Fail-Over is Enabled on a Peer-Group

To verify fast fail-over is enabled on a particular BGP neighbor, use the `show ip bgp neighbors` command. Because fast fail-over is disabled by default, it appears only if it has been enabled (shown in bold).

```
Dell#sh ip bgp neighbors
```

```
BGP neighbor is 100.100.100.100, remote AS 65517, internal link
 Member of peer-group test for session parameters
 BGP version 4, remote router ID 30.30.30.5
 BGP state ESTABLISHED, in this state for 00:19:15
 Last read 00:00:15, last write 00:00:06
 Hold time is 180, keepalive interval is 60 seconds
 Received 52 messages, 0 notifications, 0 in queue
 Sent 45 messages, 5 notifications, 0 in queue
 Received 6 updates, Sent 0 updates
 Route refresh request: received 0, sent 0
 Minimum time between advertisement runs is 5 seconds
 Minimum time before advertisements start is 0 seconds
```

```
Capabilities received from neighbor for IPv4 Unicast :
 MULTIPROTO_EXT(1)
 ROUTE_REFRESH(2)
 CISCO_ROUTE_REFRESH(128)
```

```
Capabilities advertised to neighbor for IPv4 Unicast :
 MULTIPROTO_EXT(1)
 ROUTE_REFRESH(2)
 CISCO_ROUTE_REFRESH(128)
```

#### **fail-over enabled**

```
Update source set to Loopback 0
```

```
Peer active in peer-group outbound optimization
```

```
For address family: IPv4 Unicast
BGP table version 52, neighbor version 52
4 accepted prefixes consume 16 bytes
Prefix advertised 0, denied 0, withdrawn 0
```

To verify that fast fail-over is enabled on a peer-group, use the `show ip bgp peer-group` command (shown in bold).

```
Dell#sh ip bgp peer-group
```

```
Peer-group test
fail-over enabled
BGP version 4
Minimum time between advertisement runs is 5 seconds
```

```
For address family: IPv4 Unicast
BGP neighbor is test
Number of peers in this group 1
Peer-group members (* - outbound optimized):
```

```
100.100.100.100*
Dell#
router bgp 65517
 neighbor test peer-group
 neighbor test fail-over
 neighbor test no shutdown
```

## Configuring Passive Peering

When you enable a peer-group, the software sends an OPEN message to initiate a TCP connection. If you enable passive peering for the peer group, the software does not send an OPEN message, but it responds to an OPEN message.

When a BGP neighbor connection with authentication configured is rejected by a passive peer-group, the system does not allow another passive peer-group on the same subnet to connect with the BGP neighbor. To work around this, change the BGP configuration or change the order of the peer group configuration.

You can constrain the number of passive sessions accepted by the neighbor. The `limit` keyword allows you to set the total number of sessions the neighbor will accept, between 2 and 256. The default is **256** sessions.

1. Configure a peer group that does not initiate TCP connections with other peers.

CONFIG-ROUTER-BGP mode

```
neighbor peer-group-name peer-group passive limit
```

Enter the `limit` keyword to restrict the number of sessions accepted.

2. Assign a subnet to the peer group.

CONFIG-ROUTER-BGP mode

```
neighbor peer-group-name subnet subnet-number mask
```

The peer group responds to OPEN messages sent on this subnet.

3. Enable the peer group.

CONFIG-ROUTER-BGP mode

```
neighbor peer-group-name no shutdown
```

4. Create and specify a remote peer for BGP neighbor.

CONFIG-ROUTER-BGP mode

```
neighbor peer-group-name remote-as as-number
```

Only after the peer group responds to an OPEN message sent on the subnet does its BGP state change to ESTABLISHED. After the peer group is ESTABLISHED, the peer group is the same as any other peer group.

For more information about peer groups, refer to [Configuring Peer Groups](#).

## Maintaining Existing AS Numbers During an AS Migration

The local-as feature smooths out the BGP network migration operation and allows you to maintain existing ASNs during a BGP network migration.

When you complete your migration, be sure to reconfigure your routers with the new information and disable this feature.

- Allow external routes from this neighbor.

CONFIG-ROUTERBGP mode

```
neighbor {IP address | peer-group-name local-as as number [no prepend]
- Peer Group Name: 16 characters.
- AS-number: 0 to 65535 (2-Byte) or 1 to 4294967295 (4-Byte) or 0.1 to 65535.65535 (Dotted format).
- No Prepend: specifies that local AS values are not prepended to announcements from the neighbor.
```

Format: IP Address: A.B.C.D.

You must use [Configuring Peer Groups](#) before assigning it to an AS. This feature is not supported on passive peer groups.

### Example of the Verifying that Local AS Numbering is Disabled

The first line in bold shows the actual AS number. The second two lines in bold show the local AS number (6500) maintained during migration.

To disable this feature, use the `no neighbor local-as` command in CONFIGURATION ROUTER BGP mode.

```
R2(conf-router_bgp)#show conf
!
router bgp 65123
 bgp router-id 192.168.10.2
 network 10.10.21.0/24
 network 10.10.32.0/24
 network 100.10.92.0/24
 network 192.168.10.0/24
 bgp four-octet-as-support
 neighbor 10.10.21.1 remote-as 65123
 neighbor 10.10.21.1 filter-list Laura in
 neighbor 10.10.21.1 no shutdown
 neighbor 10.10.32.3 remote-as 65123
 neighbor 10.10.32.3 no shutdown
 neighbor 100.10.92.9 remote-as 65192
 neighbor 100.10.92.9 local-as 6500
 neighbor 100.10.92.9 no shutdown
 neighbor 192.168.10.1 remote-as 65123
 neighbor 192.168.10.1 update-source Loopback 0
```

### Allowing an AS Number to Appear in its Own AS Path

This command allows you to set the number of times a particular AS number can occur in the AS path. The allow-as feature permits a BGP speaker to allow the ASN to be present for a specified number of times in the update received from the peer, even if that ASN matches its own. The AS-PATH loop is detected if the local ASN is present more than the specified number of times in the command.



- Allow this neighbor ID to use the AS path the specified number of times.

CONFIG-ROUTER-BGP mode

```
neighbor {IP address | peer-group-name} allowas-in number
```

- Peer Group Name: 16 characters.
- Number: 1 through 10.

Format: IP Address: A.B.C.D.

You must use [Configuring Peer Groups](#) before assigning it to an AS.

### Example of Viewing AS Numbers in AS Paths

The lines shown in bold are the number of times ASN 65123 can appear in the AS path (**allows-in 9**).

To disable this feature, use the `no neighbor allow-as in number` command in CONFIGURATION ROUTER BGP mode.

```
R2(conf-router_bgp)#show conf
!
router bgp 65123
 bgp router-id 192.168.10.2
 network 10.10.21.0/24
 network 10.10.32.0/24
 network 100.10.92.0/24
 network 192.168.10.0/24
 bgp four-octet-as-support
 neighbor 10.10.21.1 remote-as 65123
 neighbor 10.10.21.1 filter-list Laura in
 neighbor 10.10.21.1 no shutdown
 neighbor 10.10.32.3 remote-as 65123
 neighbor 10.10.32.3 no shutdown
 neighbor 100.10.92.9 remote-as 65192
 neighbor 100.10.92.9 local-as 6500
 neighbor 100.10.92.9 no shutdown
 neighbor 192.168.10.1 remote-as 65123
 neighbor 192.168.10.1 update-source Loopback 0
 neighbor 192.168.10.1 no shutdown
```

### Enabling Graceful Restart

Use this feature to lessen the negative effects of a BGP restart.

The Dell Networking OS advertises support for this feature to BGP neighbors through a capability advertisement. You can enable graceful restart by router and/or by peer or peer group.

 **NOTE:** By default, BGP graceful restart is disabled.

The default role for BGP is as a receiving or restarting peer. If you enable BGP, when a peer that supports graceful restart resumes operating, The Dell Networking OS performs the following tasks:

- Continues saving routes received from the peer if the peer advertised it had graceful restart capability. Continues forwarding traffic to the peer.
- Flags routes from the peer as Stale and sets a timer to delete them if the peer does not perform a graceful restart.
- Deletes all routes from the peer if forwarding state information is not saved.

- Speeds convergence by advertising a special update packet known as an end-of-RIB marker. This marker indicates the peer has been updated with all routes in the local RIB.

If you configure your system to do so, the system can perform the following actions during a hot failover:

- Save all forwarding information base (FIB) and content addressable memory (CAM) entries on the line card and continue forwarding traffic while the secondary route processor module (RPM) is coming online.
- Advertise to all BGP neighbors and peer-groups that the forwarding state of all routes has been saved. This prompts all peers to continue saving the routes they receive and to continue forwarding traffic.
- Bring the secondary RPM online as the primary and re-open sessions with all peers operating in No Shutdown mode.
- Defer best path selection for a certain amount of time. This helps optimize path selection and results in fewer updates being sent out.

To enable graceful restart, use the `configure router bgp graceful-restart` command.

- Enable graceful restart for the BGP node.

CONFIG-ROUTER-BGP mode

```
bgp graceful-restart
```

- Set maximum restart time for all peers.

CONFIG-ROUTER-BGP mode

```
bgp graceful-restart [restart-time time-in-seconds]
```

The default is **120 seconds**.

- Set maximum time to retain the restarting peer's stale paths.

CONFIG-ROUTER-BGP mode

```
bgp graceful-restart [stale-path-time time-in-seconds]
```

The default is **360 seconds**.

- Local router supports graceful restart as a receiver only.

CONFIG-ROUTER-BGP mode

```
bgp graceful-restart [role receiver-only]
```

## Enabling Neighbor Graceful Restart

BGP graceful restart is active only when the neighbor becomes established. Otherwise, it is disabled. Graceful-restart applies to all neighbors with established adjacency.

With the graceful restart feature, the system enables the receiving/restarting mode by default. In Receiver-Only mode, graceful restart saves the advertised routes of peers that support this capability when they restart. This option provides support for remote peers for their graceful restart without supporting the feature itself.

You can implement BGP graceful restart either by neighbor or by BGP peer-group. For more information, refer to the *Dell Networking OS Command Line Interface Reference Guide*.

- Add graceful restart to a BGP neighbor or peer-group.

CONFIG-ROUTER-BGP mode

- ```
neighbor {ip-address | peer-group-name} graceful-restart
```
- Set the maximum restart time for the neighbor or peer-group.
CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} graceful-restart [restart-time time-in-seconds]
```

The default is **120 seconds**.

- Local router supports graceful restart for this neighbor or peer-group as a receiver only.
CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} graceful-restart [role receiver-only]
```

- Set the maximum time to retain the restarting neighbor's or peer-group's stale paths.
CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} graceful-restart [stale-path-time time-in-seconds]
```

The default is **360 seconds**.

Filtering on an AS-Path Attribute

You can use the BGP attribute, AS_PATH, to manipulate routing policies.

The AS_PATH attribute contains a sequence of AS numbers representing the route's path. As the route traverses an AS, the ASN is prepended to the route. You can manipulate routes based on their AS_PATH to affect interdomain routing. By identifying certain ASN in the AS_PATH, you can permit or deny routes based on the number in its AS_PATH.

AS-PATH ACLs use regular expressions to search AS_PATH values. AS-PATH ACLs have an "implicit deny." This means that routes that do not meet any Match filter are dropped.

To configure an AS-PATH ACL to filter a specific AS_PATH value, use these commands in the following sequence.

1. Assign a name to a AS-PATH ACL and enter AS-PATH ACL mode.
CONFIGURATION mode

```
ip as-path access-list as-path-name
```

2. Enter the parameter to match BGP AS-PATH for filtering.
CONFIG-AS-PATH mode

```
{deny | permit} filter parameter
```

This is the filter that is used to match the AS-path. The entries can be any format, letters, numbers, or regular expressions.

You can enter this command multiple times if multiple filters are desired.

For accepted expressions, refer to [Regular Expressions as Filters](#).

3. Return to CONFIGURATION mode.
AS-PATH ACL mode

```
exit
```

4. Enter ROUTER BGP mode.

CONFIGURATION mode

```
router bgp as-number
```

5. Use a configured AS-PATH ACL for route filtering and manipulation.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} filter-list as-path-name {in | out}
```

If you assign an non-existent or empty AS-PATH ACL, the software allows all routes.

Example of the show ip bgp paths Command

To view all BGP path attributes in the BGP database, use the `show ip bgp paths` command in EXEC Privilege mode.

```
Dell#show ip bgp paths
Total 30655 Paths
Address      Hash Refcount Metric Path
0x4014154 0    3          18508 701 3549 19421 i
0x4013914 0    3          18508 701 7018 14990 i
0x5166d6c 0    3          18508 209 4637 1221 9249 9249 i
0x5e62df4 0    2          18508 701 17302 i
0x3a1814c 0    26         18508 209 22291 i
0x567ea9c 0    75         18508 209 3356 2529 i
0x6cc1294 0    2          18508 209 1239 19265 i
0x6cc18d4 0    1          18508 701 2914 4713 17935 i
0x5982e44 0    162        18508 209 i
0x67d4a14 0    2          18508 701 19878 ?
0x559972c 0    31         18508 209 18756 i
0x59cd3b4 0    2          18508 209 7018 15227 i
0x7128114 0    10         18508 209 3356 13845 i
0x536a914 0    3          18508 209 701 6347 7781 i
0x2ffe884 0    1          18508 701 3561 9116 21350 i
0x2ff7284 0    99         18508 701 1239 577 855 ?
0x2ff7ec4 0    4          18508 209 3561 4755 17426 i
0x2ff8544 0    3          18508 701 5743 2648 i
0x736c144 0    1          18508 701 209 568 721 1494 i
0x3b8d224 0    10         18508 209 701 2019 i
0x5ebl44 0    1          18508 701 8584 16158 i
0x5cd891c 0    9          18508 209 6453 4759 i
--More--
```

Regular Expressions as Filters

Regular expressions are used to filter AS paths or community lists. A regular expression is a special character used to define a pattern that is then compared with an input string.

For an AS-path access list, as shown in the previous commands, if the AS path matches the regular expression in the access list, the route matches the access list.

The following lists the regular expressions accepted in the Dell Networking OS.

Regular Expression Definition

^ (caret)	Matches the beginning of the input string. Alternatively, when used as the first character within brackets [^], this matches any number except the ones specified within the brackets.
\$ (dollar)	Matches the end of the input string.
. (period)	Matches any single character, including white space.
* (asterisk)	Matches 0 or more sequences of the immediately previous character or pattern.
+ (plus)	Matches 1 or more sequences of the immediately previous character or pattern.
? (question)	Matches 0 or 1 sequence of the immediately previous character or pattern.
() (parenthesis)	Specifies patterns for multiple use when one of the multiplier metacharacters follows: asterisk *, plus sign +, or question mark ?
[] (brackets)	Matches any enclosed character and specifies a range of single characters.
- (hyphen)	Used within brackets to specify a range of AS or community numbers.
_ (underscore)	Matches a ^, a \$, a comma, a space, or a {, or a }. Placed on either side of a string to specify a literal and disallow substring matching. You can precede or follow numerals enclosed by underscores by any of the characters listed.
 (pipe)	Matches characters on either side of the metacharacter; logical OR.

As seen in the following example, the expressions are displayed when using the `show` commands. To view the AS-PATH ACL configuration, use the `show config` command in CONFIGURATION AS-PATH ACL mode and the `show ip as-path-access-list` command in EXEC Privilege mode.

For more information about this command and route filtering, refer to [Filtering BGP Routes](#).

The following example applies access list Eagle to routes inbound from BGP peer 10.5.5.2. Access list Eagle uses a regular expression to deny routes originating in AS 32. The first lines shown in bold create the access list and filter. The second lines shown in bold are the regular expression shown as part of the access list filter.

Example of Using Regular Expression to Filter AS Paths

```
Dell(config)#router bgp 99
Dell(conf-router_bgp)#neigh AAA peer-group
Dell(conf-router_bgp)#neigh AAA no shut
Dell(conf-router_bgp)#show conf
!
router bgp 99
  neighbor AAA peer-group
  neighbor AAA no shutdown
  neighbor 10.155.15.2 remote-as 32
  neighbor 10.155.15.2 shutdown
Dell(conf-router_bgp)#neigh 10.155.15.2 filter-list 1 in
Dell(conf-router_bgp)#ex
```

```
Dell(conf)#ip as-path access-list Eagle
```

```
Dell(config-as-path)#deny 32$
```

```
Dell(config-as-path)#ex
```

```
Dell(conf)#router bgp 99
```

```
Dell(conf-router_bgp)#neighbor AAA filter-list Eagle in
```

```
Dell(conf-router_bgp)#show conf
!
router bgp 99
  neighbor AAA peer-group
  neighbor AAA filter-list Eaglein
  neighbor AAA no shutdown
  neighbor 10.155.15.2 remote-as 32
  neighbor 10.155.15.2 filter-list 1 in
  neighbor 10.155.15.2 shutdown
Dell(conf-router_bgp)#ex
```

Redistributing Routes

In addition to filtering routes, you can add routes from other routing instances or protocols to the BGP process. With the `redistribute` command, you can include ISIS, OSPF, static, or directly connected routes in the BGP process.

To add routes from other routing instances or protocols, use any of the following commands in ROUTER BGP mode.

- Include, directly connected or user-configured (static) routes in BGP.

ROUTER BGP or CONF-ROUTER_BGPv6_AF mode

```
redistribute {connected | static} [route-map map-name]
```

Configure the `map-name` parameter to specify the name of a configured route map.

- Include specific ISIS routes in BGP.

ROUTER BGP or CONF-ROUTER_BGPv6_AF mode

```
redistribute isis [level-1 | level-1-2 | level-2] [metric value] [route-map map-name]
```

Configure the following parameters:

- `level-1`, `level-1-2`, or `level-2`: Assign all redistributed routes to a level. The default is **level-2**.
- `metric value`: The value is from 0 to 16777215. The default is **0**.
- `map-name`: name of a configured route map.

- Include specific OSPF routes in IS-IS.

ROUTER BGP or CONF-ROUTER_BGPv6_AF mode


```
redistribute ospf process-id [match external {1 | 2} | match internal]
[metric-type {external | internal}] [route-map map-name]
```

Configure the following parameters:

- `process-id`: the range is from 1 to 65535.
- `match external`: the range is from 1 or 2.
- `match internal`
- `metric-type`: external or internal.
- `map-name`: name of a configured route map.

Enabling Additional Paths

The add-path feature is disabled by default.

 **NOTE:** Note: In some cases, while receiving 1K same routes from more than 64 iBGP neighbors, BGP sessions holdtime of 10 seconds may flap. The BGP add-path does not update packets for advertisement and cannot scale to higher numbers. Either reduce the number of routes added or increase the holddown timer value.

To allow multiple paths sent to peers, use the following commands.

1. Allow the advertisement of multiple paths for the same address prefix without the new paths replacing any previous ones.

CONFIG-ROUTER-BGP mode

```
bgp add-path [send | receive | both] count
```

The range is from 2 to 64.

2. Allow the specified neighbor/peer group to send/ receive multiple path advertisements.

CONFIG-ROUTER-BGP mode

```
neighbor add-path
```

3. Configure the maximum number of parallel routes (multipath support) BGP supports.

CONFIG-ROUTER-BGP mode

```
max-path number
```

The range is from 2 to 64.

Configuring IP Community Lists

Within the Dell Networking OS, you have multiple methods of manipulating routing attributes. One attribute you can manipulate is the COMMUNITY attribute. This attribute is an optional attribute that is defined for a group of destinations. You can assign a COMMUNITY attribute to BGP routers by using an IP community list. After you create an IP community list, you can apply routing decisions to all routers meeting the criteria in the IP community list.

IETF RFC 1997 defines the COMMUNITY attribute and the predefined communities of INTERNET, NO_EXPORT_SUBCONFED, NO_ADVERTISE, and NO_EXPORT. All BGP routes belong to the INTERNET community. In the RFC, the other communities are defined as follows:

- All routes with the NO_EXPORT_SUBCONFED (0xFFFFF03) community attribute are not sent to CONFED-EBGP or EBGP peers, but are sent to IBGP peers within CONFED-SUB-AS.
- All routes with the NO_ADVERTISE (0xFFFFF02) community attribute must not be advertised.
- All routes with the NO_EXPORT (0xFFFFF01) community attribute must not be advertised outside a BGP boundary, but are sent to CONFED-EBGP and IBGP peers.

The Dell Networking OS also supports BGP Extended Communities as described in RFC 4360 — BGP Extended Communities Attribute.

To configure an IP community list, use these commands.

1. Create a community list and enter COMMUNITY-LIST mode.

CONFIGURATION mode

```
ip community-list community-list-name
```

2. Configure a community list by denying or permitting specific community numbers or types of community.

CONFIG-COMMUNITYLIST mode

```
{deny | permit} {community-number | local-AS | no-advertise | no-export |  
quote-regexp regular-expression-list | regexp regular-expression}
```

- *community-number*: use AA:NN format where AA is the AS number (2 Bytes or 4 Bytes) and NN is a value specific to that autonomous system.
- *local-AS*: routes with the COMMUNITY attribute of NO_EXPORT_SUBCONFED.
- *no-advertise*: routes with the COMMUNITY attribute of NO_ADVERTISE.
- *no-export*: routes with the COMMUNITY attribute of NO_EXPORT.
- *quote-regexp*: then any number of regular expressions. The software applies all regular expressions in the list.
- *regexp*: then a regular expression.

Example of the show ip community-lists Command

To view the configuration, use the `show config` command in CONFIGURATION COMMUNITY-LIST or CONFIGURATION EXTCOMMUNITY LIST mode or the `show ip {community-lists | extcommunity-list}` command in EXEC Privilege mode.

```
Dell#show ip community-lists  
ip community-list standard 1  
deny 701:20  
deny 702:20  
deny 703:20  
deny 704:20  
deny 705:20  
deny 14551:20  
deny 701:112  
deny 702:112  
deny 703:112  
deny 704:112  
deny 705:112
```

Configuring an IP Extended Community List

To configure an IP extended community list, use these commands.

1. Create a extended community list and enter the EXTCOMMUNITY-LIST mode.

CONFIGURATION mode

```
ip extcommunity-list extcommunity-list-name
```

2. Two types of extended communities are supported.

CONFIG-COMMUNITY-LIST mode

```
{permit | deny} {{rt | soo} {ASN:NN | IPADDR:N} | regex REGEX-LINE}
```

Filter routes based on the type of extended communities they carry using one of the following keywords:

- *rt*: route target.

- `soo`: route origin or site-of-origin. Support for matching extended communities against regular expression is also supported. Match against a regular expression using the following keyword.
- `regex`: regular expression.

Example of the `show ip extcommunity-lists` Command

To set or modify an extended community attribute, use the `set extcommunity {rt | soo} {ASN:NN | IPADDR:NN}` command.

To view the configuration, use the `show config` command in CONFIGURATION COMMUNITY-LIST or CONFIGURATION EXTCOMMUNITY LIST mode or the `show ip {community-lists | extcommunity-list}` command in EXEC Privilege mode.

```
Dell#show ip community-lists
ip community-list standard 1
deny 701:20
deny 702:20
deny 703:20
deny 704:20
deny 705:20
deny 14551:20
deny 701:112
deny 702:112
deny 703:112
deny 704:112
deny 705:112
```

Filtering Routes with Community Lists

To use an IP community list or IP extended community list to filter routes, you must apply a match community filter to a route map and then apply that route map to a BGP neighbor or peer group.

1. Enter the ROUTE-MAP mode and assign a name to a route map.

CONFIGURATION mode

```
route-map map-name [permit | deny] [sequence-number]
```

2. Configure a match filter for all routes meeting the criteria in the IP community or IP extended community list.

CONFIG-ROUTE-MAP mode

```
match {community community-list-name [exact] | extcommunity extcommunity-list-name [exact]}
```

3. Return to CONFIGURATION mode.

CONFIG-ROUTE-MAP mode

```
exit
```

4. Enter ROUTER BGP mode.

CONFIGURATION mode

```
router bgp as-number
```

AS-number: 0 to 65535 (2-Byte) or 1 to 4294967295 (4-Byte) or 0.1 to 65535.65535 (Dotted format)

5. Apply the route map to the neighbor or peer group's incoming or outgoing routes.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} route-map map-name {in | out}
```

To view the BGP configuration, use the `show config` command in CONFIGURATION ROUTER BGP mode. To view a route map configuration, use the `show route-map` command in EXEC Privilege mode. To view which BGP routes meet an IP community or IP extended community list's criteria, use the `show ip bgp {community-list | extcommunity-list}` command in EXEC Privilege mode.

Manipulating the COMMUNITY Attribute

In addition to permitting or denying routes based on the values of the COMMUNITY attributes, you can manipulate the COMMUNITY attribute value and send the COMMUNITY attribute with the route information.

By default, the system does not send the COMMUNITY attribute.

To send the COMMUNITY attribute to BGP neighbors, use the following command.

- Enable the software to send the router's COMMUNITY attribute to the BGP neighbor or peer group specified.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} send-community
```

To view the BGP configuration, use the `show config` command in CONFIGURATION ROUTER BGP mode.

If you want to remove or add a specific COMMUNITY number from a BGP path, you must create a route map with one or both of the following statements in the route map. Then apply that route map to a BGP neighbor or peer group.

1. Enter ROUTE-MAP mode and assign a name to a route map.

CONFIGURATION mode

```
route-map map-name [permit | deny] [sequence-number]
```

2. Configure a set filter to delete all COMMUNITY numbers in the IP community list.

CONFIG-ROUTE-MAP mode

```
set comm-list community-list-name delete
```

OR

```
set community {community-number | local-as | no-advertise | no-export | none}
```

Configure a community list by denying or permitting specific community numbers or types of community.

- `community-number`: use AA:NN format where AA is the AS number (2 or 4 Bytes) and NN is a value specific to that autonomous system.
- `local-AS`: routes with the COMMUNITY attribute of NO_EXPORT_SUBCONFED and are not sent to EBGp peers.
- `no-advertise`: routes with the COMMUNITY attribute of NO_ADVERTISE and are not advertised.
- `no-export`: routes with the COMMUNITY attribute of NO_EXPORT.

- none: remove the COMMUNITY attribute.
 - additive: add the communities to already existing communities.
3. Return to CONFIGURATION mode.

CONFIG-ROUTE-MAP mode

exit

4. Enter the ROUTER BGP mode.

CONFIGURATION mode

router bgp *as-number*

5. Apply the route map to the neighbor or peer group's incoming or outgoing routes.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} route-map map-name {in | out}
```

Example of the show ip bgp community Command

To view the BGP configuration, use the `show config` command in CONFIGURATION ROUTER BGP mode. To view a route map configuration, use the `show route-map` command in EXEC Privilege mode.

To view BGP routes matching a certain community number or a pre-defined BGP community, use the `show ip bgp community` command in EXEC Privilege mode.

```
Dell>show ip bgp community
BGP table version is 3762622, local router ID is 10.114.8.48
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
* i 3.0.0.0/8	195.171.0.16	100	0	0	209 701 80 i
*>i 4.2.49.12/30	195.171.0.16	100	0	0	209 i
* i 4.21.132.0/23	195.171.0.16	100	0	0	209 6461 16422 i
*>i 4.24.118.16/30	195.171.0.16	100	0	0	209 i
*>i 4.24.145.0/30	195.171.0.16	100	0	0	209 i
*>i 4.24.187.12/30	195.171.0.16	100	0	0	209 i
*>i 4.24.202.0/30	195.171.0.16	100	0	0	209 i
*>i 4.25.88.0/30	195.171.0.16	100	0	0	209 3561 3908 i
*>i 6.1.0.0/16	195.171.0.16	100	0	0	209 7170 1455 i
*>i 6.2.0.0/22	195.171.0.16	100	0	0	209 7170 1455 i
*>i 6.3.0.0/18	195.171.0.16	100	0	0	209 7170 1455 i

--More--

Changing MED Attributes

By default, the system uses the MULTI_EXIT_DISC or MED attribute when comparing EBGp paths from the same AS.

To change how the MED attribute is used, enter any or all of the following commands.

- Enable MED comparison in the paths from neighbors with different ASs.

CONFIG-ROUTER-BGP mode

```
bgp always-compare-med
```

By default, this comparison is not performed.

- Change the bestpath MED selection.

CONFIG-ROUTER-BGP mode

```
bgp bestpath med {confed | missing-as-best}
```

- `confed`: Chooses the bestpath MED comparison of paths learned from BGP confederations.
- `missing-as-best`: Treat a path missing an MED as the most preferred one.

To view the nondefault values, use the `show config` command in CONFIGURATION ROUTER BGP mode.

Changing the LOCAL_PREFERENCE Attribute

In the Dell Networking OS, you can change the value of the LOCAL_PREFERENCE attribute. To change the default values of this attribute for all routes received by the router, use the following command.

- Change the LOCAL_PREF value.

CONFIG-ROUTER-BGP mode

```
bgp default local-preference value
```

- `value`: the range is from 0 to 4294967295.

The default is **100**.

To view the BGP configuration, use the `show config` command in CONFIGURATION ROUTER BGP mode or the `show running-config bgp` command in EXEC Privilege mode.

A more flexible method for manipulating the LOCAL_PREF attribute value is to use a route map.

1. Enter the ROUTE-MAP mode and assign a name to a route map.

CONFIGURATION mode

```
route-map map-name [permit | deny] [sequence-number]
```

2. Change LOCAL_PREF value for routes meeting the criteria of this route map.

CONFIG-ROUTE-MAP mode

```
set local-preference value
```

3. Return to CONFIGURATION mode.

CONFIG-ROUTE-MAP mode

```
exit
```

4. Enter ROUTER BGP mode.

CONFIGURATION mode

```
router bgp as-number
```

5. Apply the route map to the neighbor or peer group's incoming or outgoing routes.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} route-map map-name {in | out}
```

To view the BGP configuration, use the `show config` command in CONFIGURATION ROUTER BGP mode. To view a route map configuration, use the `show route-map` command in EXEC Privilege mode.

Changing the NEXT_HOP Attribute

You can change how the NEXT_HOP attribute is used.

To change how the NEXT_HOP attribute is used, enter the first command. To view the BGP configuration, use the `show config` command in CONFIGURATION ROUTER BGP mode or the `show running-config bgp` command in EXEC Privilege mode.

You can also use route maps to change this and other BGP attributes. For example, you can include the second command in a route map to specify the next hop address.

- Disable next hop processing and configure the router as the next hop for a BGP neighbor.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} next-hop-self
```

- Sets the next hop address.

CONFIG-ROUTE-MAP mode

```
set next-hop ip-address
```

Changing the WEIGHT Attribute

To change how the WEIGHT attribute is used, enter the first command. You can also use route maps to change this and other BGP attributes. For example, you can include the second command in a route map to specify the next hop address.

- Assign a weight to the neighbor connection.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} weight weight
```

- *weight*: the range is from 0 to 65535.

The default is **0**.

- Sets weight for the route.

CONFIG-ROUTE-MAP mode

```
set weight weight
```

- *weight*: the range is from 0 to 65535.

To view BGP configuration, use the `show config` command in CONFIGURATION ROUTER BGP mode or the `show running-config bgp` command in EXEC Privilege mode.

Enabling Multipath

By default, the software allows one path to a destination. You can enable multipath to allow up to 16 parallel paths to a destination.

To allow more than one path, use the following command.

The `show ip bgp network` command includes multipath information for that network.

- Enable multiple parallel paths.

CONFIG-ROUTER-BGP mode


```
maximum-paths {ebgp | ibgp} number
```


The `show ip bgp network` command includes multipath information for that network.

Filtering BGP Routes

Filtering routes allows you to implement BGP policies.

You can use either IP prefix lists, route maps, AS-PATH ACLs or IP community lists (using a route map) to control which routes the BGP neighbor or peer group accepts and advertises. Prefix lists filter routes based on route and prefix length, while AS-Path ACLs filter routes based on the ASN. Route maps can filter and set conditions, change attributes, and assign update policies.

 **NOTE:** The Dell Networking OS supports up to 255 characters in a set community statement inside a route map.


 **NOTE:** With the Dell Networking OS, you can create inbound and outbound policies. Each of the commands used for filtering has `in` and `out` parameters that you must apply. In the system, the order of preference varies depending on whether the attributes are applied for inbound updates or outbound updates.

For inbound and outbound updates the order of preference is:

- prefix lists (using the `neighbor distribute-list` command)
- AS-PATH ACLs (using the `neighbor filter-list` command)
- route maps (using the `neighbor route-map` command)

Prior to filtering BGP routes, create the prefix list, AS-PATH ACL, or route map.

For configuration information about prefix lists, AS-PATH ACLs, and route maps, refer to [Access Control Lists \(ACLs\)](#).

 **NOTE:** When you configure a new set of BGP policies, to ensure the changes are made, always reset the neighbor or peer group by using the `clear ip bgp` command in EXEC Privilege mode.

To filter routes using prefix lists, use the following commands.

1. Create a prefix list and assign it a name.

CONFIGURATION mode

```
ip prefix-list prefix-name
```

2. Create multiple prefix list filters with a deny or permit action.

CONFIG-PREFIX LIST mode

```
seq sequence-number {deny | permit} {any | ip-prefix [ge | le] }
```

- `ge`: minimum prefix length to be matched.
- `le`: maximum prefix length to be matched.

For information about configuring prefix lists, refer to [Access Control Lists \(ACLs\)](#).

3. Return to CONFIGURATION mode.

CONFIG-PREFIX LIST mode

```
exit
```

4. Enter ROUTER BGP mode.

CONFIGURATION mode

```
router bgp as-number
```

5. Filter routes based on the criteria in the configured prefix list.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} distribute-list prefix-list-name {in  
| out}
```

Configure the following parameters:

- *ip-address* or *peer-group-name*: enter the neighbor's IP address or the peer group's name.
- *prefix-list-name*: enter the name of a configured prefix list.
- *in*: apply the prefix list to inbound routes.
- *out*: apply the prefix list to outbound routes.

As a reminder, the following are rules concerning prefix lists:

- If the prefix list contains no filters, all routes are permitted.
- If none of the routes match any of the filters in the prefix list, the route is denied. This action is called an implicit deny. (If you want to forward all routes that do not match the prefix list criteria, you must configure a prefix list filter to permit all routes. For example, you could have the following filter as the last filter in your prefix list `permit 0.0.0.0/0 le 32`).
- After a route matches a filter, the filter's action is applied. No additional filters are applied to the route.

To view the BGP configuration, use the `show config` command in ROUTER BGP mode. To view a prefix list configuration, use the `show ip prefix-list detail` or `show ip prefix-list summary` commands in EXEC Privilege mode.

Filtering BGP Routes Using Route Maps

To filter routes using a route map, use these commands.

1. Create a route map and assign it a name.

CONFIGURATION mode

```
route-map map-name [permit | deny] [sequence-number]
```

2. Create multiple route map filters with a match or set action.

CONFIG-ROUTE-MAP mode

```
{match | set}
```

For information about configuring route maps, refer to [Access Control Lists \(ACLs\)](#).

3. Return to CONFIGURATION mode.

CONFIG-ROUTE-MAP mode

```
exit
```

4. Enter ROUTER BGP mode.

CONFIGURATION mode

```
router bgp as-number
```

5. Filter routes based on the criteria in the configured route map.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} route-map map-name {in | out}
```

Configure the following parameters:

- *ip-address* or *peer-group-name*: enter the neighbor's IP address or the peer group's name.
- *map-name*: enter the name of a configured route map.
- *in*: apply the route map to inbound routes.
- *out*: apply the route map to outbound routes.

To view the BGP configuration, use the `show config` command in CONFIGURATION ROUTER BGP mode. To view a route map configuration, use the `show route-map` command in EXEC Privilege mode.

Filtering BGP Routes Using AS-PATH Information

To filter routes based on AS-PATH information, use these commands.

1. Create a AS-PATH ACL and assign it a name.

CONFIGURATION mode

```
ip as-path access-list as-path-name
```

2. Create a AS-PATH ACL filter with a deny or permit action.

AS-PATH ACL mode

```
{deny | permit} as-regular-expression
```

3. Return to CONFIGURATION mode.

AS-PATH ACL

```
exit
```

4. Enter ROUTER BGP mode.

CONFIGURATION mode

```
router bgp as-number
```

5. Filter routes based on the criteria in the configured route map.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} filter-list as-path-name {in | out}
```

Configure the following parameters:

- *ip-address* or *peer-group-name*: enter the neighbor's IP address or the peer group's name.
- *as-path-name*: enter the name of a configured AS-PATH ACL.
- *in*: apply the AS-PATH ACL map to inbound routes.
- *out*: apply the AS-PATH ACL to outbound routes.

To view which commands are configured, use the `show config` command in CONFIGURATION ROUTER BGP mode and the `show ip as-path-access-list` command in EXEC Privilege mode.

To forward all routes not meeting the AS-PATH ACL criteria, include the **permit .*** filter in your AS-PATH ACL.

Configuring BGP Route Reflectors

BGP route reflectors are intended for ASs with a large mesh; they reduce the amount of BGP control traffic.

With route reflection configured properly, IBGP routers are not fully meshed within a cluster but all receive routing information.

Configure clusters of routers where one router is a concentration router and the others are clients who receive their updates from the concentration router.

To configure a route reflector, use the following commands.

- Assign an ID to a router reflector cluster.

CONFIG-ROUTER-BGP mode

```
bgp cluster-id cluster-id
```

You can have multiple clusters in an AS.

- Configure the local router as a route reflector and the neighbor or peer group identified is the route reflector client.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} route-reflector-client
```

When you enable a route reflector, the system automatically enables route reflection to all clients. To disable route reflection between all clients in this reflector, use the `no bgp client-to-client reflection` command in CONFIGURATION ROUTER BGP mode. All clients must be fully meshed before you disable route reflection.

To view a route reflector configuration, use the `show config` command in CONFIGURATION ROUTER BGP mode or the `show running-config bgp` in EXEC Privilege mode.

Aggregating Routes

The Dell Networking OS provides multiple ways to aggregate routes in the BGP routing table. At least one specific route of the aggregate must be in the routing table for the configured aggregate to become active.

To aggregate routes, use the following command.

AS_SET includes AS_PATH and community information from the routes included in the aggregated route.

- Assign the IP address and mask of the prefix to be aggregated.

CONFIG-ROUTER-BGP mode

```
aggregate-address ip-address mask [advertise-map map-name] [as-set]  
[attribute-map map-name] [summary-only] [suppress-map map-name]
```

Example of Viewing Aggregated Routes

In the `show ip bgp` command, aggregates contain an 'a' in the first column (shown in bold) and routes suppressed by the aggregate contain an 's' in the first column.

```
Dell#show ip bgp  
BGP table version is 0, local router ID is 10.101.15.13  
Status codes: s suppressed, d damped, h history, * valid, > best  
Path source: I - internal, a - aggregate, c - confed-external, r -
```

redistributed, n - network
Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 7.0.0.0/29	10.114.8.33	0		0 18508	?
*> 7.0.0.0/30	10.114.8.33	0		0 18508	?
*>a 9.0.0.0/8	192.0.0.0			32768 18508	701 {7018 2686 3786} ?

Configuring BGP Confederations

Another way to organize routers within an AS and reduce the mesh for IBGP peers is to configure BGP confederations.

As with route reflectors, BGP confederations are recommended only for IBGP peering involving many IBGP peering sessions per router. Basically, when you configure BGP confederations, you break the AS into smaller sub-AS, and to those outside your network, the confederations appear as one AS. Within the confederation sub-AS, the IBGP neighbors are fully meshed and the MED, NEXT_HOP, and LOCAL_PREF attributes are maintained between confederations.

To configure BGP confederations, use the following commands.

- Specifies the confederation ID.

CONFIG-ROUTER-BGP mode

```
bgp confederation identifier as-number
```

– *as-number*: from 0 to 65535 (2 Byte) or from 1 to 4294967295 (4 Byte).

- Specifies which confederation sub-AS are peers.

CONFIG-ROUTER-BGP mode

```
bgp confederation peers as-number [... as-number]
```

– *as-number*: from 0 to 65535 (2 Byte) or from 1 to 4294967295 (4 Byte).

All Confederation routers must be either 4 Byte or 2 Byte. You cannot have a mix of router ASN support.

To view the configuration, use the `show config` command in CONFIGURATION ROUTER BGP mode.

Enabling Route Flap Dampening

When EBGp routes become unavailable, they “flap” and the router issues both WITHDRAWN and UPDATE notices.

A flap is when a route:

- is withdrawn
- is readvertised after being withdrawn
- has an attribute change

The constant router reaction to the WITHDRAWN and UPDATE notices causes instability in the BGP process. To minimize this instability, you may configure penalties (a numeric value) for routes that flap. When that penalty value reaches a configured limit, the route is not advertised, even if the route is up. In Dell, that penalty value is 1024. As time passes and the route does not flap, the penalty value decrements or is decayed. However, if the route flaps again, it is assigned another penalty.

The penalty value is cumulative and penalty is added under following cases:

- Withdraw
- Readvertise
- Attribute change

When dampening is applied to a route, its path is described by one of the following terms:

- history entry — an entry that stores information on a downed route
- dampened path — a path that is no longer advertised
- penalized path — a path that is assigned a penalty

To configure route flap dampening parameters, set dampening parameters using a route map, clear information on route dampening and return suppressed routes to active state, view statistics on route flapping, or change the path selection from the default mode (deterministic) to non-deterministic, use the following commands.

- Enable route dampening.
CONFIG-ROUTER-BGP mode

```
bgp dampening [half-life | reuse | suppress max-suppress-time] [route-map
map-name]
```

Enter the following optional parameters to configure route dampening parameters:

- *half-life*: the range is from 1 to 45. Number of minutes after which the Penalty is decreased. After the router assigns a Penalty of 1024 to a route, the Penalty is decreased by half after the half-life period expires. The default is **15 minutes**.
 - *reuse*: the range is from 1 to 20000. This number is compared to the flapping route's Penalty value. If the Penalty value is less than the reuse value, the flapping route is once again advertised (or no longer suppressed). Withdrawn routes are removed from history state. The default is **750**.
 - *suppress*: the range is from 1 to 20000. This number is compared to the flapping route's Penalty value. If the Penalty value is greater than the suppress value, the flapping route is no longer advertised (that is, it is suppressed). The default is **2000**.)
 - *max-suppress-time*: the range is from 1 to 255. The maximum number of minutes a route can be suppressed. The default is four times the half-life value. The default is **60 minutes**.
 - *route-map map-name*: name of a configured route map. Only match commands in the configured route map are supported. Use this parameter to apply route dampening to selective routes.
- Enter the following optional parameters to configure route dampening.
CONFIG-ROUTE-MAP mode

```
set dampening half-life reuse suppress max-suppress-time
```

- *half-life*: the range is from 1 to 45. Number of minutes after which the Penalty is decreased. After the router assigns a Penalty of 1024 to a route, the Penalty is decreased by half after the half-life period expires. The default is **15 minutes**.
 - *reuse*: the range is from 1 to 20000. This number is compared to the flapping route's Penalty value. If the Penalty value is less than the reuse value, the flapping route is once again advertised (or no longer suppressed). The default is **750**.
 - *suppress*: the range is from 1 to 20000. This number is compared to the flapping route's Penalty value. If the Penalty value is greater than the suppress value, the flapping route is no longer advertised (that is, it is suppressed). The default is **2000**.
 - *max-suppress-time*: the range is from 1 to 255. The maximum number of minutes a route can be suppressed. The default is four times the half-life value. The default is **60 minutes**.
- Clear all information or only information on a specific route.

EXEC Privilege

```
clear ip bgp dampening [ip-address mask]
```

- View all flap statistics or for specific routes meeting the following criteria.

EXEC or EXEC Privilege mode

```
show ip bgp flap-statistics [ip-address [mask]] [filter-list as-path-name]
[regexp regular-expression]
```

- `ip-address [mask]`: enter the IP address and mask.
- `filter-list as-path-name`: enter the name of an AS-PATH ACL.
- `regexp regular-expression`: enter a regular express to match on.

By default, the path selection in Dell is deterministic, that is, paths are compared irrespective of the order of their arrival. You can change the path selection method to non-deterministic, that is, paths are compared in the order in which they arrived (starting with the most recent). Furthermore, in non-deterministic mode, the software may not compare MED attributes though the paths are from the same AS.

- Change the best path selection method to non-deterministic.

CONFIG-ROUTER-BGP mode

```
bgp non-deterministic-med
```



NOTE: When you change the best path selection method, path selection for existing paths remains unchanged until you reset it by entering the `clear ip bgp` command in EXEC Privilege mode.

Example of Configuring a Route for Reuse or Restart

Example of Viewing the Number of Dampened Routes

To view the BGP configuration, use the `show config` command in CONFIGURATION ROUTER BGP mode or the `show running-config bgp` command in EXEC Privilege mode.

The following example shows how to configure values to reuse or restart a route. In the following example, `default = 15` is the set time before the value decrements, `bgp dampening 2 ?` is the set re-advertise value, `bgp dampening 2 2000 ?` is the suppress value, and `bgp dampening 2 2000 3000 ?` is the time to suppress a route. Default values are also shown.

```
Dell(conf-router_bgp)#bgp dampening ?
<1-45> Half-life time for the penalty (default = 15)
route-map Route-map to specify criteria for dampening
<cr>
Dell(conf-router_bgp)#bgp dampening 2 ?
<1-20000> Value to start reusing a route (default = 750)
Dell(conf-router_bgp)#bgp dampening 2 2000 ?
<1-20000> Value to start suppressing a route (default = 2000)
Dell(conf-router_bgp)#bgp dampening 2 2000 3000 ?
<1-255> Maximum duration to suppress a stable route (default = 60)
Dell(conf-router_bgp)#bgp dampening 2 2000 3000 10 ?
route-map Route-map to specify criteria for dampening
<cr>
```

To view a count of dampened routes, history routes, and penalized routes when you enable route dampening, look at the seventh line of the `show ip bgp summary` command output, as shown in the following example (bold).

```
Dell>show ip bgp summary
BGP router identifier 10.114.8.131, local AS number 65515
BGP table version is 855562, main routing table version 780266
122836 network entrie(s) and 221664 paths using 29697640 bytes of memory
34298 BGP path attribute entrie(s) using 1920688 bytes of memory
29577 BGP AS-PATH entrie(s) using 1384403 bytes of memory
184 BGP community entrie(s) using 7616 bytes of memory
Dampening enabled. 0 history paths, 0 dampened paths, 0 penalized paths

Neighbor      AS      MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
10.114.8.34   18508   82883   79977   780266  0    2 00:38:51  118904
10.114.8.33   18508   117265  25069   780266  0    20 00:38:50  102759
Dell>
```

To view which routes are dampened (non-active), use the `show ip bgp dampened-routes` command in EXEC Privilege mode.

Changing BGP Timers

To configure BGP timers, use either or both of the following commands.

Timer values configured with the `neighbor timers` command override the timer values configured with the `timers bgp` command.

When two neighbors, configured with different `keepalive` and `holdtime` values, negotiate for new values, the resulting values are as follows:

- the lower of the `holdtime` values is the new `holdtime` value, and
- whichever is the lower value; one-third of the new `holdtime` value, or the configured `keepalive` value is the new `keepalive` value.

- Configure timer values for a BGP neighbor or peer group.

CONFIG-ROUTER-BGP mode

```
neighbors {ip-address | peer-group-name} timers keepalive holdtime
```

- `keepalive`: the range is from 1 to 65535. Time interval, in seconds, between keepalive messages sent to the neighbor routers. The default is **60 seconds**.
- `holdtime`: the range is from 3 to 65536. Time interval, in seconds, between the last keepalive message and declaring the router dead. The default is **180 seconds**.

- Configure timer values for all neighbors.

CONFIG-ROUTER-BGP mode

```
timers bgp keepalive holdtime
```

- `keepalive`: the range is from 1 to 65535. Time interval, in seconds, between keepalive messages sent to the neighbor routers. The default is **60 seconds**.
- `holdtime`: the range is from 3 to 65536. Time interval, in seconds, between the last keepalive message and declaring the router dead. The default is **180 seconds**.

To view non-default values, use the `show config` command in CONFIGURATION ROUTER BGP mode or the `show running-config bgp` command in EXEC Privilege mode.

Enabling BGP Neighbor Soft-Reconfiguration

BGP soft-reconfiguration allows for faster and easier route changing.

Changing routing policies typically requires a reset of BGP sessions (the TCP connection) for the policies to take effect. Such resets cause undue interruption to traffic due to hard reset of the BGP cache and the time it takes to re-establish the session. BGP soft reconfig allows for policies to be applied to a session without clearing the BGP Session. Soft-reconfig can be done on a per-neighbor basis and can either be inbound or outbound.

BGP soft-reconfiguration clears the policies without resetting the TCP connection.

To reset a BGP connection using BGP soft reconfiguration, use the `clear ip bgp` command in EXEC Privilege mode at the system prompt.

When you enable soft-reconfiguration for a neighbor and you execute the `clear ip bgp soft in` command, the update database stored in the router is replayed and updates are reevaluated. With this command, the replay and update process is triggered only if a route-refresh request is not negotiated with the peer. If the request is indeed negotiated (after execution of `clear ip bgp soft in`), BGP sends a route-refresh request to the neighbor and receives all of the peer's updates.

To use soft reconfiguration (or soft reset) without preconfiguration, both BGP peers must support the soft route refresh capability, which is advertised in the open message sent when the peers establish a TCP session.

To determine whether a BGP router supports this capability, use the `show ip bgp neighbors` command. If a router supports the route refresh capability, the following message displays: `Received route refresh capability from peer.`

If you specify a BGP peer group by using the `peer-group-name` argument, all members of the peer group inherit the characteristic configured with this command.

- Clear all information or only specific details.

EXEC Privilege mode

```
clear ip bgp {* | neighbor-address | AS Numbers | ipv4 | peer-group-name}
[soft [in | out]]
```

- *: Clears all peers.
- neighbor-address: Clears the neighbor with this IP address.
- AS Numbers: Peers' AS numbers to be cleared.
- ipv4: Clears information for the IPv4 address family.
- peer-group-name: Clears all members of the specified peer group.

- Enable soft-reconfiguration for the BGP neighbor specified.

CONFIG-ROUTER-BGP mode

```
neighbor {ip-address | peer-group-name} soft-reconfiguration inbound
```

BGP stores all the updates received by the neighbor but does not reset the peer-session.

Entering this command starts the storage of updates, which is required to do inbound soft reconfiguration. Outbound BGP soft reconfiguration does not require inbound soft reconfiguration to be enabled.

Example of Soft-Reconfiguration of a BGP Neighbor

The example enables inbound soft reconfiguration for the neighbor 10.108.1.1. All updates received from this neighbor are stored unmodified, regardless of the inbound policy. When inbound soft reconfiguration is done later, the stored information is used to generate a new set of inbound updates.

```
Dell>router bgp 100
  neighbor 10.108.1.1 remote-as 200
  neighbor 10.108.1.1 soft-reconfiguration inbound
```

Route Map Continue

The BGP route map continue feature, `continue [sequence-number]`, (in ROUTE-MAP mode) allows movement from one route-map entry to a specific route-map entry (the sequence number).

If you do not specify a sequence number, the continue feature moves to the next sequence number (also known as an “implied continue”). If a match clause exists, the continue feature executes only after a successful match occurs. If there are no successful matches, continue is ignored.

Match a Clause with a Continue Clause

The continue feature can exist without a match clause.

Without a match clause, the continue clause executes and jumps to the specified route-map entry. With a match clause and a continue clause, the match clause executes first and the continue clause next in a specified route map entry. The continue clause launches only after a successful match. The behavior is:

- A successful match with a continue clause—the route map executes the set clauses and then goes to the specified route map entry after execution of the continue clause.
- If the next route map entry contains a continue clause, the route map executes the continue clause if a successful match occurs.
- If the next route map entry does not contain a continue clause, the route map evaluates normally. If a match does not occur, the route map does not continue and falls-through to the next sequence number, if one exists

Set a Clause with a Continue Clause

If the route-map entry contains sets with the continue clause, the set actions operation is performed first followed by the continue clause jump to the specified route map entry.

- If a set actions operation occurs in the first route map entry and then the same set action occurs with a different value in a subsequent route map entry, the last set of actions overrides the previous set of actions with the same `set` command.
- If the `set community additive` and `set as-path prepend` commands are configured, the communities and AS numbers are prepended.

Enabling MBGP Configurations

Multiprotocol BGP (MBGP) is an enhanced BGP that carries IP multicast routes. BGP carries two sets of routes: one set for unicast routing and one set for multicast routing. The routes associated with multicast routing are used by the protocol independent multicast (PIM) to build data distribution trees.

The Dell Networking OS MBGP is implemented per RFC 1858. You can enable the MBGP feature per router and/or per peer/peer-group.

The default is **IPv4 Unicast** routes.

When you configure a peer to support IPv4 multicast, the system takes the following actions:

- Send a capability advertisement to the peer in the BGP Open message specifying IPv4 multicast as a supported AFI/SAFI (Subsequent Address Family Identifier).
- If the corresponding capability is received in the peer's Open message, BGP marks the peer as supporting the AFI/SAFI.
- When exchanging updates with the peer, BGP sends and receives IPv4 multicast routes if the peer is marked as supporting that AFI/SAFI.
- Exchange of IPv4 multicast route information occurs through the use of two new attributes called MP_REACH_NLRI and MP_UNREACH_NLRI, for feasible and withdrawn routes, respectively.
- If the peer has not been activated in any AFI/SAFI, the peer remains in Idle state.

Most Dell Networking OS BGP IPv4 unicast commands are extended to support the IPv4 multicast RIB using extra options to the command. For a detailed description of the MBGP commands, refer to the *Dell Networking OS Command Line Interface Reference Guide*.

- Enables support for the IPv4 multicast family on the BGP node.

CONFIG-ROUTER-BGP mode

```
address family ipv4 multicast
```

- Enable IPv4 multicast support on a BGP neighbor/peer group.

CONFIG-ROUTER-BGP-AF (Address Family) mode

```
neighbor [ip-address | peer-group-name] activate
```

BGP Regular Expression Optimization

The Dell Networking OS optimizes processing time when using regular expressions by caching and re-using regular expression evaluated results, at the expense of some memory in RP1 processor.

BGP policies that contain regular expressions to match against as-paths and communities might take a lot of CPU processing time, thus affect BGP routing convergence. Also, `show bgp` commands that get filtered through regular expressions can take a lot of CPU cycles, especially when the database is large.

This feature is turned on by default. If necessary, use the `bgp regex-eval-optz-disable` command in CONFIGURATION ROUTER BGP mode to disable it.

Debugging BGP

To enable BGP debugging, use any of the following commands.

- View all information about BGP, including BGP events, keepalives, notifications, and updates.

EXEC Privilege mode

```
debug ip bgp [ip-address | peer-group peer-group-name] [in | out]
```

- View information about BGP route being dampened.

EXEC Privilege mode

```
debug ip bgp dampening [in | out]
```


- View information about local BGP state changes and other BGP events.

EXEC Privilege mode

```
debug ip bgp [ip-address | peer-group peer-group-name] events [in | out]
```

- View information about BGP KEEPALIVE messages.

EXEC Privilege mode

```
debug ip bgp [ip-address | peer-group peer-group-name] keepalive [in | out]
```

- View information about BGP notifications received from or sent to neighbors.

EXEC Privilege mode

```
debug ip bgp [ip-address | peer-group peer-group-name] notifications [in | out]
```

- View information about BGP updates and filter by prefix name.

EXEC Privilege mode

```
debug ip bgp [ip-address | peer-group peer-group-name] updates [in | out]
[prefix-list name]
```

- Enable soft-reconfiguration debug.

EXEC Privilege mode

```
debug ip bgp {ip-address | peer-group-name} soft-reconfiguration
```

In-BGP is shown using the `show ip protocols` command.

The system displays debug messages on the console. To view which debugging commands are enabled, use the `show debugging` command in EXEC Privilege mode.

To disable a specific debug command, use the keyword `no` then the debug command. For example, to disable debugging of BGP updates, use `no debug ip bgp updates` command.

To disable all BGP debugging, use the `no debug ip bgp` command.

To disable all debugging, use the `undebug all` command.

Storing Last and Bad PDUs

The system stores the last notification sent/received and the last bad protocol data unit (PDU) received on a per peer basis. The last bad PDU is the one that causes a notification to be issued.

In the following example, the last seven lines shown in bold are the last PDUs.

Example of the `show ip bgp neighbor` Command to View Last and Bad PDUs

```
Dell(conf-router_bgp)#do show ip bgp neighbors 1.1.1.2
```

```
BGP neighbor is 1.1.1.2, remote AS 2, external link
BGP version 4, remote router ID 2.4.0.1
BGP state ESTABLISHED, in this state for 00:00:01
Last read 00:00:00, last write 00:00:01
Hold time is 90, keepalive interval is 30 seconds
Received 1404 messages, 0 in queue
  3 opens, 1 notifications, 1394 updates
  6 keepalives, 0 route refresh requests
```

```

Sent 48 messages, 0 in queue
  3 opens, 2 notifications, 0 updates
  43 keepalives, 0 route refresh requests
Minimum time between advertisement runs is 30 seconds
Minimum time before advertisements start is 0 seconds

Capabilities received from neighbor for IPv4 Unicast :
MULTIPROTO_EXT(1)
ROUTE_REFRESH(2)
CISCO_ROUTE_REFRESH(128)

Capabilities advertised to neighbor for IPv4 Unicast :
MULTIPROTO_EXT(1)
ROUTE_REFRESH(2)
CISCO_ROUTE_REFRESH(128)
For address family: IPv4 Unicast
BGP table version 1395, neighbor version 1394
Prefixes accepted 1 (consume 4 bytes), 0 withdrawn by peer
Prefixes advertised 0, rejected 0, 0 withdrawn from peer

Connections established 3; dropped 2
Last reset 00:00:12, due to Missing well known attribute

Notification History
'UPDATE error/Missing well-known attr' Sent : 1 Recv: 0
'Connection Reset' Sent : 1 Recv: 0

```

PDU Counters

The Dell Networking OS version 7.5.1.0 introduces additional counters for various types of PDUs sent and received from neighbors.

These are seen in the output of the `show ip bgp neighbor` command.

Sample Configurations

The following example configurations show how to enable BGP and set up some peer groups. These examples are not comprehensive directions. They are intended to give you some guidance with typical configurations.

To support your own IP addresses, interfaces, names, and so on, you can copy and paste from these examples to your CLI. Be sure that you make the necessary changes.

The following illustration shows the configurations described on the following examples. These configurations show how to create BGP areas using physical and virtual links. They include setting up the interfaces and peers groups with each other.

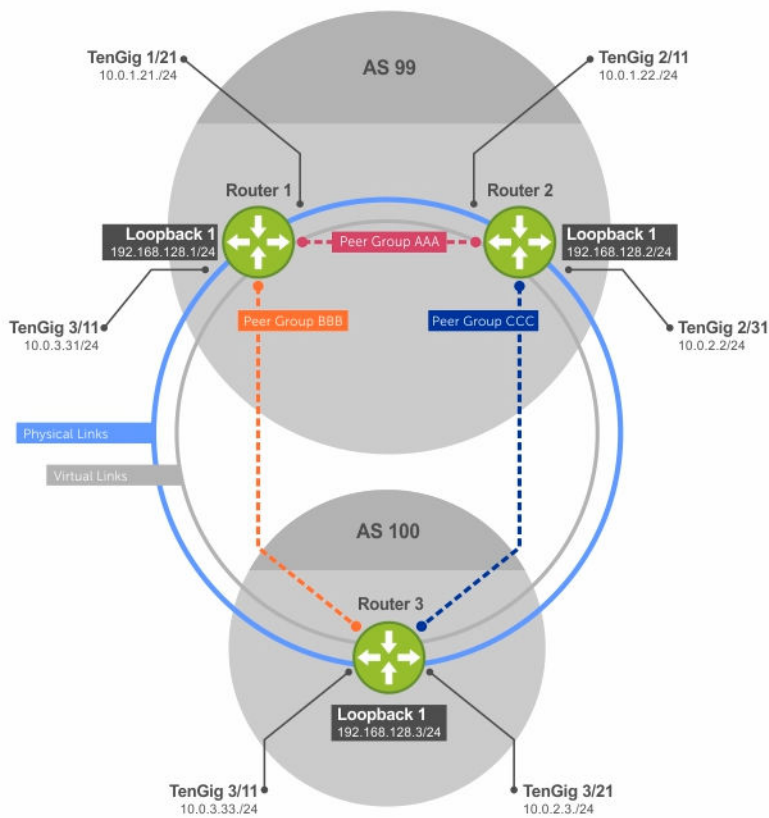


Figure 24. Sample Configurations

Example of Enabling BGP (Router 1)

Example of Enabling BGP (Router 2)

Example of Enabling BGP (Router 3)

Example of Enabling Peer Groups (Router 1)

Example of Enabling Peer Groups (Router 2)

Example of Enabling Peer Groups (Router 3)

```
R1# conf
R1(conf)#int loop 0
R1(conf-if-lo-0)#ip address 192.168.128.1/24
R1(conf-if-lo-0)#no shutdown
R1(conf-if-lo-0)#show config
!
  interface Loopback 0
  ip address 192.168.128.1/24
no shutdown
R1(conf-if-lo-0)#int gig 1/21
R1(conf-if-gi-1/21)#ip address 10.0.1.21/24
R1(conf-if-gi-1/21)#no shutdown
R1(conf-if-gi-1/21)#show config
!
  interface GigabitEthernet 1/21
  ip address 10.0.1.21/24
```

```

no shutdown
R1(conf-if-gi-1/21)#int gig 1/31
R1(conf-if-gi-1/31)#ip address 10.0.3.31/24
R1(conf-if-gi-1/31)#no shutdown
R1(conf-if-gi-1/31)#show config
!
    interface GigabitEthernet 1/31
    ip address 10.0.3.31/24
no shutdown
R1(conf-if-gi-1/31)#router bgp 99
R1(conf-router_bgp)#network 192.168.128.0/24
R1(conf-router_bgp)#neighbor 192.168.128.2 remote 99
R1(conf-router_bgp)#neighbor 192.168.128.2 no shut
R1(conf-router_bgp)#neighbor 192.168.128.2 update-source loop 0
R1(conf-router_bgp)#neighbor 192.168.128.3 remote 100
R1(conf-router_bgp)#neighbor 192.168.128.3 no shut
R1(conf-router_bgp)#neighbor 192.168.128.3 update-source loop 0
R1(conf-router_bgp)#show config
!
router bgp 99
    network 192.168.128.0/24
    neighbor 192.168.128.2 remote-as 99
    neighbor 192.168.128.2 update-source Loopback 0
    neighbor 192.168.128.2 no shutdown
    neighbor 192.168.128.3 remote-as 100
    neighbor 192.168.128.3 update-source Loopback 0
    neighbor 192.168.128.3 no shutdown
R1(conf-router_bgp)#end
R1#
R1#show ip bgp summary
BGP router identifier 192.168.128.1, local AS number 99
BGP table version is 4, main routing table version 4
4 network entrie(s) using 648 bytes of memory
6 paths using 408 bytes of memory
BGP-RIB over all using 414 bytes of memory
3 BGP path attribute entrie(s) using 144 bytes of memory
2 BGP AS-PATH entrie(s) using 74 bytes of memory
2 neighbor(s) using 8672 bytes of memory

Neighbor      AS  MsgRcvd  MsgSent  TblVer  InQ  OutQ  Up/Down  State/Pfx
192.168.128.2 99   4         5         4        0    0    00:00:32  1
192.168.128.3 100  5         4         1        0    0    00:00:09  4
R1#

R2# conf
R2(conf)#int loop 0
R2(conf-if-lo-0)#ip address 192.168.128.2/24
R2(conf-if-lo-0)#no shutdown
R2(conf-if-lo-0)#show config
!
    interface Loopback 0
    ip address 192.168.128.2/24
no shutdown
R2(conf-if-lo-0)#int gig 2/11
R2(conf-if-gi-2/11)#ip address 10.0.1.22/24
R2(conf-if-gi-2/11)#no shutdown
R2(conf-if-gi-2/11)#show config
!
interface GigabitEthernet 2/11
    ip address 10.0.1.22/24
    no shutdown
R2(conf-if-gi-2/11)#int gig 2/31

```

```

R2(conf-if-gi-2/31)#ip address 10.0.2.2/24
R2(conf-if-gi-2/31)#no shutdown
R2(conf-if-gi-2/31)#show config
!
interface GigabitEthernet 2/31
ip address 10.0.2.2/24
no shutdown
R2(conf-if-gi-2/31)#
R2(conf-if-gi-2/31)#router bgp 99
R2(conf-router_bgp)#network 192.168.128.0/24
R2(conf-router_bgp)#neighbor 192.168.128.1 remote 99
R2(conf-router_bgp)#neighbor 192.168.128.1 no shut
R2(conf-router_bgp)#neighbor 192.168.128.1 update-source loop 0
R2(conf-router_bgp)#neighbor 192.168.128.3 remote 100
R2(conf-router_bgp)#neighbor 192.168.128.3 no shut
R2(conf-router_bgp)#neighbor 192.168.128.3 update loop 0
R2(conf-router_bgp)#show config
!
router bgp 99
  bgp router-id 192.168.128.2
  network 192.168.128.0/24
  bgp graceful-restart
  neighbor 192.168.128.1 remote-as 99
  neighbor 192.168.128.1 update-source Loopback 0
  neighbor 192.168.128.1 no shutdown
  neighbor 192.168.128.3 remote-as 100
  neighbor 192.168.128.3 update-source Loopback 0
  neighbor 192.168.128.3 no shutdown
R2(conf-router_bgp)#end

R2#show ip bgp summary
BGP router identifier 192.168.128.2, local AS number 99
BGP table version is 1, main routing table version 1
1 network entrie(s) using 132 bytes of memory
3 paths using 204 bytes of memory
BGP-RIB over all using 207 bytes of memory
2 BGP path attribute entrie(s) using 128 bytes of memory
2 BGP AS-PATH entrie(s) using 90 bytes of memory
2 neighbor(s) using 9216 bytes of memory

Neighbor      AS  MsgRcvd  MsgSent  TblVer  InQ  OutQ  Up/Down  State/Pfx
192.168.128.1 99   40       35       1        0    0    00:01:05  1
192.168.128.3 100  4        4        1        0    0    00:00:16  1
R2#

R3# conf
R3(conf)#
R3(conf)#int loop 0
R3(conf-if-lo-0)#ip address 192.168.128.3/24
R3(conf-if-lo-0)#no shutdown
R3(conf-if-lo-0)#show config
!
  interface Loopback 0
  ip address 192.168.128.3/24
no shutdown
R3(conf-if-lo-0)#int gig 3/11
R3(conf-if-gi-3/11)#ip address 10.0.3.33/24
R3(conf-if-gi-3/11)#no shutdown
R3(conf-if-gi-3/11)#show config
!
interface GigabitEthernet 3/11
  ip address 10.0.3.33/24
  no shutdown

```

```

R3(conf-if-lo-0)#int gig 3/21
R3(conf-if-gi-3/21)#ip address 10.0.2.3/24
R3(conf-if-gi-3/21)#no shutdown
R3(conf-if-gi-3/21)#show config
!
interface GigabitEthernet 3/21
 ip address 10.0.2.3/24
 no shutdown

R3(conf-if-gi-3/21)#
R3(conf-if-gi-3/21)#router bgp 100
R3(conf-router_bgp)#show config
!
router bgp 100
R3(conf-router_bgp)#network 192.168.128.0/24
R3(conf-router_bgp)#neighbor 192.168.128.1 remote 99
R3(conf-router_bgp)#neighbor 192.168.128.1 no shut
R3(conf-router_bgp)#neighbor 192.168.128.1 update-source loop 0
R3(conf-router_bgp)#neighbor 192.168.128.2 remote 99
R3(conf-router_bgp)#neighbor 192.168.128.2 no shut
R3(conf-router_bgp)#neighbor 192.168.128.2 update loop 0
R3(conf-router_bgp)#show config
!
router bgp 100
 network 192.168.128.0/24
 neighbor 192.168.128.1 remote-as 99
 neighbor 192.168.128.1 update-source Loopback 0
 neighbor 192.168.128.1 no shutdown
 neighbor 192.168.128.2 remote-as 99
 neighbor 192.168.128.2 update-source Loopback 0
 neighbor 192.168.128.2 no shutdown
R3(conf)#end
R3#show ip bgp summary
BGP router identifier 192.168.128.3, local AS number 100
BGP table version is 1, main routing table version 1
1 network entrie(s) using 132 bytes of memory
3 paths using 204 bytes of memory
BGP-RIB over all using 207 bytes of memory
2 BGP path attribute entrie(s) using 128 bytes of memory
2 BGP AS-PATH entrie(s) using 90 bytes of memory
2 neighbor(s) using 9216 bytes of memory
Neighbor      AS  MsgRcvd  MsgSent  TblVer  InQ  OutQ  Up/Down  State/Pfx
192.168.128.1 99   24       25       1       0    0     00:14:20  1
192.168.128.2 99   14       14       1       0    0     00:10:22  1
R3#

```

```

R1#conf
R1(conf)#router bgp 99
R1(conf-router_bgp)# network 192.168.128.0/24
R1(conf-router_bgp)# neighbor AAA peer-group
R1(conf-router_bgp)# neighbor AAA no shutdown
R1(conf-router_bgp)# neighbor BBB peer-group
R1(conf-router_bgp)# neighbor BBB no shutdown
R1(conf-router_bgp)# neighbor 192.168.128.2 peer-group AAA
R1(conf-router_bgp)# neighbor 192.168.128.3 peer-group BBB
R1(conf-router_bgp)#
R1(conf-router_bgp)#show config
!
router bgp 99
 network 192.168.128.0/24
 neighbor AAA peer-group
 neighbor AAA no shutdown

```

```

neighbor BBB peer-group
neighbor BBB no shutdown
neighbor 192.168.128.2 remote-as 99
neighbor 192.168.128.2 peer-group AAA
neighbor 192.168.128.2 update-source Loopback 0
neighbor 192.168.128.2 no shutdown
neighbor 192.168.128.3 remote-as 100
neighbor 192.168.128.3 peer-group BBB
neighbor 192.168.128.3 update-source Loopback 0
neighbor 192.168.128.3 no shutdown
R1#
R1#show ip bgp summary
BGP router identifier 192.168.128.1, local AS number 99
BGP table version is 1, main routing table version 1
1 network entrie(s) using 132 bytes of memory
3 paths using 204 bytes of memory
BGP-RIB over all using 207 bytes of memory
2 BGP path attribute entrie(s) using 96 bytes of memory
2 BGP AS-PATH entrie(s) using 74 bytes of memory
2 neighbor(s) using 8672 bytes of memory

Neighbor      AS  MsgRcvd  MsgSent  TblVer  InQ  OutQ  Up/Down  State/Pfx
192.168.128.2 99  23       24       1       0   (0)   00:00:17  1
192.168.128.3 100 30       29       1       0   (0)   00:00:14  1
!
R1#show ip bgp neighbors

BGP neighbor is 192.168.128.2, remote AS 99, internal link
Member of peer-group AAA for session parameters
BGP version 4, remote router ID 192.168.128.2
BGP state ESTABLISHED, in this state for 00:00:37
Last read 00:00:36, last write 00:00:36
Hold time is 180, keepalive interval is 60 seconds
Received 23 messages, 0 in queue
  2 opens, 0 notifications, 2 updates
  19 keepalives, 0 route refresh requests
Sent 24 messages, 0 in queue
  2 opens, 1 notifications, 2 updates
  19 keepalives, 0 route refresh requests
Minimum time between advertisement runs is 5 seconds
Minimum time before advertisements start is 0 seconds
Capabilities received from neighbor for IPv4 Unicast :
  MULTIPROTO_EXT(1)
  ROUTE_REFRESH(2)
  CISCO_ROUTE_REFRESH(128)
Capabilities advertised to neighbor for IPv4 Unicast :
  MULTIPROTO_EXT(1)
  ROUTE_REFRESH(2)
  CISCO_ROUTE_REFRESH(128)

Update source set to Loopback 0
Peer active in peer-group outbound optimization
For address family: IPv4 Unicast
BGP table version 1, neighbor version 1
Prefixes accepted 1 (consume 4 bytes), withdrawn 0 by peer
Prefixes advertised 1, denied 0, withdrawn 0 from peer

Connections established 2; dropped 1
Last reset 00:00:57, due to user reset

Notification History
'Connection Reset' Sent : 1 Recv: 0
Last notification (len 21) sent 00:00:57 ago
  ffffffff ffffffff ffffffff ffffffff 00150306 00000000

```

```

Local host: 192.168.128.1, Local port: 179
Foreign host: 192.168.128.2, Foreign port: 65464
BGP neighbor is 192.168.128.3, remote AS 100, external link
  Member of peer-group BBB for session parameters
  BGP version 4, remote router ID 192.168.128.3
  BGP state ESTABLISHED, in this state for 00:00:37
  Last read 00:00:36, last write 00:00:36
  Hold time is 180, keepalive interval is 60 seconds
  Received 30 messages, 0 in queue
    4 opens, 2 notifications, 4 updates
    20 keepalives, 0 route refresh requests
  Sent 29 messages, 0 in queue
    4 opens, 1 notifications, 4 updates
    20 keepalives, 0 route refresh requests
  Minimum time between advertisement runs is 30 seconds
  Minimum time before advertisements start is 0 seconds

  Capabilities received from neighbor for IPv4 Unicast :
    MULTIPROTO_EXT(1)
Capabilities received from neighbor for IPv4 Unicast :
  MULTIPROTO_EXT(1)
  ROUTE_REFRESH(2)
  CISCO_ROUTE_REFRESH(128)
Update source set to Loopback 0
Peer active in peer-group outbound optimization
For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Prefixes accepted 1 (consume 4 bytes), withdrawn 0 by peer
  Prefixes advertised 1, denied 0, withdrawn 0 from peer
Connections established 4; dropped 3
  Last reset 00:00:54, due to user reset
R1#

```

```

R2#conf
R2(conf)#router bgp 99
R2(conf-router_bgp)# neighbor CCC peer-group
R2(conf-router_bgp)# neighbor CC no shutdown
R2(conf-router_bgp)# neighbor BBB peer-group
R2(conf-router_bgp)# neighbor BBB no shutdown
R2(conf-router_bgp)# neighbor 192.168.128.1 peer AAA
R2(conf-router_bgp)# neighbor 192.168.128.1 no shut
R2(conf-router_bgp)# neighbor 192.168.128.3 peer BBB
R2(conf-router_bgp)# neighbor 192.168.128.3 no shut
R2(conf-router_bgp)#show conf
!
router bgp 99
  network 192.168.128.0/24
  neighbor AAA peer-group
  neighbor AAA no shutdown
  neighbor BBB peer-group
  neighbor BBB no shutdown
  neighbor 192.168.128.1 remote-as 99
  neighbor 192.168.128.1 peer-group CCC
  neighbor 192.168.128.1 update-source Loopback 0
  neighbor 192.168.128.1 no shutdown
  neighbor 192.168.128.3 remote-as 100
  neighbor 192.168.128.3 peer-group BBB
  neighbor 192.168.128.3 update-source Loopback 0
  neighbor 192.168.128.3 no shutdown
R2(conf-router_bgp)#end

```

```

R2#
R2#show ip bgp summary

```



```

BGP router identifier 192.168.128.2, local AS number 99
BGP table version is 2, main routing table version 2
1 network entrie(s) using 132 bytes of memory
3 paths using 204 bytes of memory
BGP-RIB over all using 207 bytes of memory
2 BGP path attribute entrie(s) using 128 bytes of memory
2 BGP AS-PATH entrie(s) using 90 bytes of memory
2 neighbor(s) using 9216 bytes of memory

Neighbor      AS  MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/Pfx
192.168.128.1 99  140     136     2        0    (0)   00:11:24  1
192.168.128.3 100 138     140     2        0    (0)   00:18:31  1

```

```
R2#show ip bgp neighbor
```

```

BGP neighbor is 192.168.128.1, remote AS 99, internal link
Member of peer-group AAA for session parameters
BGP version 4, remote router ID 192.168.128.1
BGP state ESTABLISHED, in this state for 00:11:42
Last read 00:00:38, last write 00:00:38
Hold time is 180, keepalive interval is 60 seconds
Received 140 messages, 0 in queue
  6 opens, 2 notifications, 19 updates
  113 keepalives, 0 route refresh requests
Sent 136 messages, 0 in queue
  12 opens, 3 notifications, 6 updates
  115 keepalives, 0 route refresh requests
Minimum time between advertisement runs is 5 seconds
Minimum time before advertisements start is 0 seconds

```

```
R3#conf
```

```

R3(conf)#router bgp 100
R3(conf-router_bgp)# neighbor AAA peer-group
R3(conf-router_bgp)# neighbor AAA no shutdown
R3(conf-router_bgp)# neighbor CCC peer-group
R3(conf-router_bgp)# neighbor CCC no shutdown
R3(conf-router_bgp)# neighbor 192.168.128.2 peer-group BBB
R3(conf-router_bgp)# neighbor 192.168.128.2 no shutdown
R3(conf-router_bgp)# neighbor 192.168.128.1 peer-group BBB
R3(conf-router_bgp)# neighbor 192.168.128.1 no shutdown
R3(conf-router_bgp)#

```

```
R3(conf-router_bgp)#end
```

```

R3#show ip bgp summary
BGP router identifier 192.168.128.3, local AS number 100
BGP table version is 1, main routing table version 1
1 network entrie(s) using 132 bytes of memory
3 paths using 204 bytes of memory
BGP-RIB over all using 207 bytes of memory
2 BGP path attribute entrie(s) using 128 bytes of memory
2 BGP AS-PATH entrie(s) using 90 bytes of memory
2 neighbor(s) using 9216 bytes of memory

```

```

Neighbor      AS  MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/Pfx
192.168.128.1 99  93      99       1        0    (0)   00:00:15  1
192.168.128.2 99  122     120      1        0    (0)   00:00:11  1

```

```
R3#show ip bgp neighbor
```

```

BGP neighbor is 192.168.128.1, remote AS 99, external link
Member of peer-group BBB for session parameters
BGP version 4, remote router ID 192.168.128.1

```

```

BGP state ESTABLISHED, in this state for 00:00:21
Last read 00:00:09, last write 00:00:08
Hold time is 180, keepalive interval is 60 seconds
Received 93 messages, 0 in queue
  5 opens, 0 notifications, 5 updates
  83 keepalives, 0 route refresh requests
Sent 99 messages, 0 in queue
  5 opens, 4 notifications, 5 updates
  85 keepalives, 0 route refresh requests
Minimum time between advertisement runs is 30 seconds
Minimum time before advertisements start is 0 seconds

Capabilities received from neighbor for IPv4 Unicast :
MULTIPROTO_EXT(1)
ROUTE_REFRESH(2)
CISCO_ROUTE_REFRESH(128)

Capabilities advertised to neighbor for IPv4 Unicast :
MULTIPROTO_EXT(1)
ROUTE_REFRESH(2)
CISCO_ROUTE_REFRESH(128)

Update source set to Loopback 0
Peer active in peer-group outbound optimization

For address family: IPv4 Unicast
BGP table version 1, neighbor version 1
Prefixes accepted 1 (consume 4 bytes), withdrawn 0 by peer
Prefixes advertised 1, denied 0, withdrawn 0 from peer

Capabilities received from neighbor for IPv4 Unicast :
MULTIPROTO_EXT(1)
ROUTE_REFRESH(2)
CISCO_ROUTE_REFRESH(128)

Capabilities advertised to neighbor for IPv4 Unicast :
MULTIPROTO_EXT(1)
ROUTE_REFRESH(2)
CISCO_ROUTE_REFRESH(128)

Update source set to Loopback 0
Peer active in peer-group outbound optimization
For address family: IPv4 Unicast
BGP table version 2, neighbor version 2
Prefixes accepted 1 (consume 4 bytes), withdrawn 0 by peer
Prefixes advertised 1, denied 0, withdrawn 0 from peer

Connections established 6; dropped 5
Last reset 00:12:01, due to Closed by neighbor

Notification History
'HOLD error/Timer expired' Sent : 1 Recv: 0
'Connection Reset' Sent : 2 Recv: 2

Last notification (len 21) received 00:12:01 ago
ffffffff ffffffff ffffffff ffffffff 00150306 00000000
Local host: 192.168.128.2, Local port: 65464
Foreign host: 192.168.128.1, Foreign port: 179

BGP neighbor is 192.168.128.3, remote AS 100, external link
Member of peer-group BBB for session parameters
BGP version 4, remote router ID 192.168.128.3
BGP state ESTABLISHED, in this state for 00:18:51
Last read 00:00:45, last write 00:00:44

```

```
Hold time is 180, keepalive interval is 60 seconds
Received 138 messages, 0 in queue
  7 opens, 2 notifications, 7 updates
  122 keepalives, 0 route refresh requests
Sent 140 messages, 0 in queue
  7 opens, 4 notifications, 7 updates
  122 keepalives, 0 route refresh requests
Minimum time between advertisement runs is 30 seconds
Minimum time before advertisements start is 0 seconds
Capabilities advertised to neighbor for IPv4 Unicast :
  MULTIPROTO_EXT(1)
Capabilities received from neighbor for IPv4 Unicast :
  MULTIPROTO_EXT(1)
  ROUTE_REFRESH(2)
  CISCO_ROUTE_REFRESH(128)
ROUTE_REFRESH(2)
  CISCO_ROUTE_REFRESH(128)

Update source set to Loopback 0
Peer active in peer-group outbound optimization
For address family: IPv4 Unicast
BGP table version 2, neighbor version 2
Prefixes accepted 1 (consume 4 bytes), withdrawn 0 by peer
Prefixes advertised 1, denied 0, withdrawn 0 from peer
```

Content Addressable Memory (CAM)


Content addressable memory (CAM) is a type of memory that stores information in the form of a lookup table.

On Dell Networking systems, CAM stores Layer 2 and Layer 3 forwarding information, access-lists (ACLs), flows, and routing policies.

CAM Allocation

Allocate space for IPV4 ACLs and quality of service (QoS) regions by using the `cam-acl` command in CONFIGURATION mode.

The CAM space is allotted in filter processor (FP) blocks. The total space allocated must equal 13 FP blocks.

 **NOTE:** There are 16 FP blocks, but the system flow requires three blocks that cannot be reallocated.

The following table lists the default CAM allocation settings.

Table 8. Default Cam Allocation Settings

CAM Allocation	Setting
L3 ACL (ipv4acl)	2
L2 ACL (l2acl)	2
IPv6 L3 ACL (ipv6acl)	0
L3 QoS (ipv4qos)	2
L2 QoS (l2qos)	0
L2PT (l2pt)	0
MAC ACLs (ipmacacl)	0
ECFMACL (ecfmacl)	0
nlbclusteracl	2
FCOEACL (fcoeacl)	4
ISCSIOPTACL (iscsioptacl)	2
VMAN QoS (vman-qos)	0
VMAN Dual QoS (vman-dual-qos)	0

The `ipv6acl` and `vman-dual-qos` allocations must be entered as a factor of 2 (2, 4, 6, 8, 10). All other profile allocations can use either even or odd numbered ranges.

You must save the new CAM settings to the startup-config (`write-mem` or `copy run start`) then reload the system for the new settings to take effect.

1. Select a `cam-acl` action.

CONFIGURATION mode

```
cam-acl [default | l2acl]
```



NOTE: Selecting default resets the CAM entries to the default settings. Select `l2acl` to allocate space for the ACLs and QoS regions.

2. Enter the number of FP blocks for each region.

EXEC Privilege mode

```
l2acl number ipv4acl number ipv6acl number, ipv4qos number l2qos number,  
l2pt number ipmacacl number ecfmacacl number nlbcluster number[vman-qos |  
vman-dual-qos number
```

3. Reload the system.

EXEC Privilege mode

```
reload
```

4. Verify that the new settings will be written to the CAM on the next boot.

EXEC Privilege mode

```
show cam-acl
```

Test CAM Usage

This command applies to both IPv4 CAM profiles, but is best used when verifying QoS optimization for IPv6 ACLs.

Use this command to determine whether sufficient ACL CAM space is available to enable a service-policy. Create a Class Map with all required ACL rules, then execute the `test cam-usage` command in Privilege mode to verify the actual CAM space required. The Status column in the command output indicates whether or not the policy can be enabled.

Example of the `test cam-usage` Command

```
Dell#test cam-usage service-policy input pmap stack-unit all
```

```
Stack-Unit | Portpipe | CAM Partition | Available CAM | Estimated CAM per Port  
| Status
```

```
-----  
-----  
2 | 0 | L2ACL | 28 | 1  
| Allowed (28)
```

View CAM-ACL Settings

View the current cam-acl settings using the `show cam-acl` command.

Example of Viewing CAM-ACL Settings

```
Dell#show cam-acl

-- Chassis Cam ACL --
      Current Settings(in block sizes)
L2Acl      : 6
Ipv4Acl    : 2
Ipv6Acl    : 0
Ipv4Qos    : 2
L2Qos      : 1
L2PT       : 0
IpMacAcl   : 0
VmanQos    : 0
VmanDualQos : 0
EcfmAcl    : 0
FcoeAcl    : 0
iscsiOptAcl : 2

-- Stack unit 5 --
      Current Settings(in block sizes)
L2Acl      : 6
Ipv4Acl    : 2
Ipv6Acl    : 0
Ipv4Qos    : 2
L2Qos      : 1
L2PT       : 0
IpMacAcl   : 0
VmanQos    : 0
VmanDualQos : 0
EcfmAcl    : 0
FcoeAcl    : 0
iscsiOptAcl : 2
Dell#
```

CAM Optimization

When you enable this command, if a Policy Map containing classification rules (ACL and/or dscp/ ip-precedence rules) is applied to more than one physical interface on the same port-pipe, only a single copy of the policy is written (only 1 FP entry is used).

When you disable this command, the system behaves as described in this chapter. However, enabling CAM optimization would apply a single rate policy FP entry. If the input service policy maps applied to several ports are the same, rate policing is applied to all the ports as a group and not individually.




Figure 26. CoPP Implemented Versus CoPP Not Implemented

Configure Control Plane Policing

The MXL switch can process maximum of 4200 PPS (packets per second). Protocols that share a single queue may experience flaps if one of the protocols receives a high rate of control traffic even though Per Protocol CoPP is applied. This happens because Queue-Based Rate Limiting is applied first.

For example, border gateway protocol (BGP) and internet control message protocol (ICMP) share same queue (Q6); Q6 has 400 PPS of bandwidth by default. The desired rate of ICMP is 100 PPS and the remaining 300 PPS is assigned to BGP. If ICMP packets come at 400 PPS, BGP packets may be dropped though ICMP packets are rate-limited to 100 PPS. You can solve this by increasing Q6 bandwidth to 700 PPS to allow both ICMP and BGP packets and then applying per-flow CoPP for ICMP and BGP packets. The setting of this Q6 bandwidth is dependent on the incoming traffic for the set of protocols sharing the

same queue. If you are not aware of the incoming protocol traffic rate, you cannot set the required queue rate limit value. You must complete queue bandwidth tuning carefully because the system cannot open up to handle any rate, including traffic coming at the line rate.

CoPP policies are assigned on a per-protocol or a per-queue basis, and are assigned in CONTROL-PLANE mode to each port-pipe.

CoPP policies are configured by creating extended ACL rules and specifying rate-limits through QoS policies. The ACLs and QoS policies are assigned as service-policies.

Configuring CoPP for Protocols

This section lists the commands necessary to create and enable the service-policies for CoPP. For complete information about creating ACLs and QoS rules, refer to [Access Control Lists \(ACLs\)](#) Access Control Lists (ACLs) and [Quality of Service \(QoS\)](#).

The basics for creating a CoPP service policy are to create a Layer 2, Layer 3, and/or an IPv6 ACL rule for the desired protocol type. Then, create a QoS input policy to rate-limit the protocol traffics according to the ACL. The ACL and QoS policies are finally assigned to a control-plane service policy for each port-pipe.

1. Create a Layer 2 extended ACL for control-plane traffic policing for a particular protocol.

CONFIGURATION mode

```
mac access-list extended name cpu-qos permit {arp | frrp | gvrp | isis |  
lacp | lldp | stp}
```

2. Create a Layer 3 extended ACL for control-plane traffic policing for a particular protocol.

CONFIGURATION mode

```
ip access-list extended name cpu-qos permit {bgp | dhcp | dhcp-relay | ftp |  
icmp | igmp | msdp | ntp | ospf | pim | ip | ssh | telnet | vrrp}
```

3. Create an IPv6 ACL for control-plane traffic policing for a particular protocol.

CONFIGURATION mode

```
ipv6 access-list name cpu-qos permit {bgp | icmp | vrrp}
```

4. Create a QoS input policy for the router and assign the policing.

CONFIGURATION mode

```
qos-policy-input name cpu-qos rate-police
```

5. Create a QoS class map to differentiate the control-plane traffic and assign to an ACL.

CONFIGURATION mode

```
class-map match-any name cpu-qos match {ip | mac | ipv6} access-group name
```

6. Create a QoS input policy map to match to the class-map and qos-policy for each desired protocol.

CONFIGURATION mode

```
policy-map-input name cpu-qos class-map name qos-policy name
```

7. Enter Control Plane mode.

CONFIGURATION mode

```
control-plane-cpuqos
```

- Assign the protocol based the service policy on the control plane. Enabling this command on a port-pipe automatically enables the ACL and QoS rules creates with the `cpu-qos` keyword.

CONTROL-PLANE mode

```
service-policy rate-limit-protocols
```

Example of Creating the IP/IPv6/MAC Extended ACL

Example of Creating the QoS Input Policy

Example of Creating the QoS Class Map

Example of Matching the QoS Class Map to the QoS Policy

Example of Creating the Control Plane Service Policy

```
Dell(conf)#ip access-list extended ospf cpu-qos
Dell(conf-ip-acl-cpuqos)#permit ospf
Dell(conf-ip-acl-cpuqos)#exit
```

```
Dell(conf)#ip access-list extended bgp cpu-qos
Dell(conf-ip-acl-cpuqos)#permit bgp
Dell(conf-ip-acl-cpuqos)#exit
```

```
Dell(conf)#mac access-list extended lacp cpu-qos
Dell(conf-mac-acl-cpuqos)#permit lacp
Dell(conf-mac-acl-cpuqos)#exit
```

```
Dell(conf)#ipv6 access-list ipv6-icmp cpu-qos
Dell(conf-ipv6-acl-cpuqos)#permit icmp
Dell(conf-ipv6-acl-cpuqos)#exit
```

```
Dell(conf)#ipv6 access-list ipv6-vrrp cpu-qos
Dell(conf-ipv6-acl-cpuqos)#permit vrrp
Dell(conf-ipv6-acl-cpuqos)#exit
```

```
Dell(conf)#qos-policy-in rate_limit_200k cpu-qos
Dell(conf-in-qos-policy-cpuqos)#rate-police 200 40 peak 500 40
Dell(conf-in-qos-policy-cpuqos)#exit
```

```
Dell(conf)#qos-policy-in rate_limit_400k cpu-qos
Dell(conf-in-qos-policy-cpuqos)#rate-police 400 50 peak 600 50
Dell(conf-in-qos-policy-cpuqos)#exit
```

```
Dell(conf)#qos-policy-in rate_limit_500k cpu-qos
Dell(conf-in-qos-policy-cpuqos)#rate-police 500 50 peak 1000 50
Dell(conf-in-qos-policy-cpuqos)#exit
```

```
Dell(conf)#class-map match-any class_ospf cpu-qos
Dell(conf-class-map-cpuqos)#match ip access-group ospf
Dell(conf-class-map-cpuqos)#exit
```

```
Dell(conf)#class-map match-any class_bgp cpu-qos
Dell(conf-class-map-cpuqos)#match ip access-group bgp
Dell(conf-class-map-cpuqos)#exit
```

```
Dell(conf)#class-map match-any class_lacp cpu-qos
Dell(conf-class-map-cpuqos)#match mac access-group lacp
Dell(conf-class-map-cpuqos)#exit
```

```
Dell(conf)#class-map match-any class-ipv6-icmp cpu-qos
Dell(conf-class-map-cpuqos)#match ipv6 access-group ipv6-icmp
Dell(conf-class-map-cpuqos)#exit
```

```
Dell(conf)#policy-map-input egressFP_rate_policy cpu-qos
Dell(conf-policy-map-in-cpuqos)#class-map class_ospf qos-policy rate_limit_500k
Dell(conf-policy-map-in-cpuqos)#class-map class_bgp qos-policy rate_limit_400k
Dell(conf-policy-map-in-cpuqos)#class-map class_lacp qos-policy rate_limit_200k
Dell(conf-policy-map-in-cpuqos)#class-map class-ipv6 qos-policy rate_limit_200k
Dell(conf-policy-map-in-cpuqos)#exit
```

```
Dell(conf)#control-plane-cpuqos
Dell(conf-control-cpuqos)#service-policy rate-limit-protocols
egressFP_rate_policy
Dell(conf-control-cpuqos)#exit
```

Configuring CoPP for CPU Queues

Controlling traffic on the CPU queues does not require ACL rules, but does require QoS policies. CoPP for CPU queues converts the input rate from kbps to pps, assuming 64 bytes is the average packet size, and applies that rate to the corresponding queue. Consequently, 1 kbps is roughly equivalent to 2 pps.

The basics for creating a CoPP service policy is to create QoS policies for the desired CPU bound queue and associate it with a particular rate-limit. The QoS policies are assigned to a control-plane service policy for each port-pipe.

1. Create a QoS input policy for the router and assign the policing.

CONFIGURATION mode

```
qos-policy-input name cpu-qos
```

2. Create an input policy-map to assign the QoS policy to the desired service queues.

CONFIGURATION mode

```
policy-map--input name cpu-qos service-queue 0 qos-policy name
```

3. Enter Control Plane mode.

CONFIGURATION mode

```
control-plane-cpuqos
```

4. Assign a CPU queue-based service policy on the control plane in cpu-qos mode. Enabling this command sets the queue rates according to those configured.

CONTROL-PLANE mode

```
service-policy rate-limit-cpu-queues name
```

Example of Creating the QoS Policy

Example of Assigning the QoS Policy to the Queues

Example of Creating the Control Plane Service Policy

```
Dell#conf
Dell(conf)#qos-policy-input cpuq_1
Dell(conf-qos-policy-in)#rate-police 3000 40 peak 500 40
Dell(conf-qos-policy-in)#exit
```

```
Dell(conf)#qos-policy-input cpuq_2
Dell(conf-qos-policy-in)#rate-police 5000 80 peak 600 50
Dell(conf-qos-policy-in)#exit
```

```

Dell(conf)#policy-map-input cpuq_rate_policy cpu-qos
Dell(conf-qos-policy-in)#service-queue 5 qos-policy cpuq_1
Dell(conf-qos-policy-in)#service-queue 6 qos-policy cpuq_2
Dell(conf-qos-policy-in)#service-queue 7 qos-policy cpuq_1

Dell#conf
Dell(conf)#control-plane
Dell(conf-control-plane)#service-policy rate-limit-cpu-queues cpuq_rate_policy

```

Show Commands

The following section describes the CoPP show commands.

To view the rates for each queue, use the `show cpu-queue rate cp` command.

Example of Viewing Queue Rates

Example of Viewing Queue Mapping

Example of Viewing Queue Mapping for MAC Protocols

Example of Viewing Queue Mapping for IPv6 Protocols

```

Dell#show cpu-queue rate cp
Service-Queue Rate (PPS)
-----
Q0             1300
Q1             300
Q2             300
Q3             300
Q4             2000
Q5             400
Q6             400
Q7             1100
Dell#

```

To view the queue mapping for each configured protocol, use the `show ip protocol-queue-mapping` command.

```

Dell#show ip protocol-queue-mapping
Protocol      Src-Port  Dst-Port  TcpFlag  Queue  EgPort  Rate (kbps)
-----
TCP (BGP)     any/179   179/any   -         Q6     CP       100
UDP (DHCP)    67/68     68/67     -         Q6/Q5  CP       -
UDP (DHCP-R)  67        67        -         Q6     CP       -
TCP (FTP)     any       21        -         Q6     CP       -
ICMP          any       any       -         Q6     CP       -
IGMP         any       any       -         Q7     CP       -
TCP (MSDP)    any/639   639/any   -         Q6     CP       -
UDP (NTP)     any       123       -         Q6     CP       -
OSPF         any       any       -         Q7     CP       -
PIM          any       any       -         Q7     CP       -
UDP (RIP)     any       520       -         Q7     CP       -
TCP (SSH)     any       22        -         Q6     CP       -
TCP (TELNET)  any       23        -         Q6     CP       -
VRRP         any       any       -         Q7     CP       -
Dell#

```

To view the queue mapping for the MAC protocols, use the `show mac protocol-queue-mapping` command.

```
Dell#show mac protocol-queue-mapping
Protocol Destination Mac      EtherType Queue EgPort Rate (kbps)
-----
ARP      any                        0x0806   Q5/Q6   CP      -
FRRP     01:01:e8:00:00:10/11     any      Q7      CP      -
LACP     01:80:c2:00:00:02       0x8809   Q7      CP      -
LLDP     any                        0x88cc   Q7      CP      -
GVRP     01:80:c2:00:00:21       any      Q7      CP      -
STP      01:80:c2:00:00:00       any      Q7      CP      -
ISIS     01:80:c2:00:00:14/15   any      Q7      CP      -
         09:00:2b:00:00:04/05   any      Q7      CP      -
```

Dell#

To view the queue mapping for IPv6 protocols, use the `show ipv6 protocol-queue-mapping` command.

```
Dell#show ipv6 protocol-queue-mapping
Protocol  Src-Port Dst-Port TcpFlag Queue EgPort Rate (kbps)
-----
TCP (BGP) any/179  179/any  -       Q6    CP    -
ICMP     any      any      -       Q6    CP    -
VRRP     any      any      -       Q7    CP    -
```

Dell#

Data Center Bridging (DCB)

Data center bridging (DCB) is supported on the FC Flex IO module installed in the MXL 10/40GbE Switch.

Ethernet Enhancements in Data Center Bridging

The following section describes DCB.

- The device supports the following DCB features:
 - Data center bridging exchange protocol (DCBx)
 - Priority-based flow control (PFC)
 - Enhanced transmission selection (ETS)

DCB refers to a set of IEEE Ethernet enhancements that provide data centers with a single, robust, converged network to support multiple traffic types, including local area network (LAN), server, and storage traffic. Through network consolidation, DCB results in reduced operational cost, simplified management, and easy scalability by avoiding the need to deploy separate application-specific networks.

For example, instead of deploying an Ethernet network for LAN traffic, include additional storage area networks (SANs) to ensure lossless Fibre Channel traffic, and a separate InfiniBand network for high-performance inter-processor computing within server clusters, only one DCB-enabled network is required in a data center. The Dell Networking switches that support a unified fabric and consolidate multiple network infrastructures use a single input/output (I/O) device called a converged network adapter (CNA).

A CNA is a computer input/output device that combines the functionality of a host bus adapter (HBA) with a network interface controller (NIC). Multiple adapters on different devices for several traffic types are no longer required.

Data center bridging satisfies the needs of the following types of data center traffic in a unified fabric:

- | | |
|------------------------|---|
| LAN traffic | LAN traffic consists of many flows that are insensitive to latency requirements, while certain applications, such as streaming video, are more sensitive to latency. Ethernet functions as a best-effort network that may drop packets in the case of network congestion. IP networks rely on transport protocols (for example, TCP) for reliable data transmission with the associated cost of greater processing overhead and performance impact. |
| Storage traffic | Storage traffic based on Fibre Channel media uses the SCSI protocol for data transfer. This traffic typically consists of large data packets with a payload of 2K bytes that cannot recover from frame loss. To successfully transport storage traffic, data center Ethernet must provide no-drop service with lossless links. |

InterProcess Communication (IPC) traffic InterProcess Communication (IPC) traffic within high-performance computing clusters to share information. Server traffic is extremely sensitive to latency requirements.

To ensure lossless delivery and latency-sensitive scheduling of storage and service traffic and I/O convergence of LAN, storage, and server traffic over a unified fabric, IEEE data center bridging adds the following extensions to a classical Ethernet network:

- 802.1Qbb — Priority-based Flow Control (PFC)
- 802.1Qaz — Enhanced Transmission Selection (ETS)
- 802.1Qau — Congestion Notification
- Data Center Bridging Exchange (DCBx) protocol

NOTE: In the Dell Networking OS version 8.3.12.0, only the PFC, ETS, and DCBx features are supported in data center bridging.

Priority-Based Flow Control

In a data center network, priority-based flow control (PFC) manages large bursts of one traffic type in multiprotocol links so that it does not affect other traffic types and no frames are lost due to congestion.

When PFC detects congestion on a queue for a specified priority, it sends a pause frame for the 802.1p priority traffic to the transmitting device. In this way, PFC ensures that PFC-enabled priority traffic is not dropped by the switch.

PFC enhances the existing 802.3x pause and 802.1p priority capabilities to enable flow control based on 802.1p priorities (classes of service). Instead of stopping all traffic on a link (as performed by the traditional Ethernet pause mechanism), PFC pauses traffic on a link according to the 802.1p priority set on a traffic type. You can create lossless flows for storage and server traffic while allowing for loss in case of LAN traffic congestion on the same physical interface.

The following illustration shows how PFC handles traffic congestion by pausing the transmission of incoming traffic with dot1p priority 3.




Figure 27. Priority-Based Flow Control

In the system, PFC is implemented as follows:

- PFC is supported on specified 802.1p priority traffic (dot1p 0 to 7) and is configured per interface. However, only two lossless queues are supported on an interface: one for Fibre Channel over

Ethernet (FCoE) converged traffic and one for Internet Small Computer System Interface (iSCSI) storage traffic. Configure the same lossless queues on all ports.

- PFC delay constraints place an upper limit on the transmit time of a queue after receiving a message to pause a specified priority.
- By default, PFC is enabled on an interface with no dot1p priorities configured. You can configure the PFC priorities if the switch negotiates with a remote peer using DCBX.
- During DCBX negotiation with a remote peer:
 - If the negotiation succeeds and the port is in DCBX Willing mode to receive a peer configuration, PFC parameters from the peer are used to configure PFC priorities on the port. If you enable the link-level flow control mechanism on the interface, DCBX negotiation with a peer is not performed.
 - If the negotiation fails and PFC is enabled on the port, any user-configured PFC input policies are applied. If no PFC input policy has been previously applied, the PFC default setting is used (no priorities configured). If you do not enable PFC on an interface, you can enable the 802.3x link-level pause function. By default, the link-level pause is disabled.
- PFC supports buffering to receive data that continues to arrive on an interface while the remote system reacts to the PFC operation.
- PFC uses the DCB MIB IEEE802.1azd2.5 and the PFC MIB IEEE802.1bb-d2.2.

Enhanced Transmission Selection

Enhanced transmission selection (ETS) supports optimized bandwidth allocation between traffic types in multiprotocol (Ethernet, FCoE, SCSI) links.

ETS allows you to divide traffic according to its 802.1p priority into different priority groups (traffic classes) and configure bandwidth allocation and queue scheduling for each group to ensure that each traffic type is correctly prioritized and receives its required bandwidth. For example, you can prioritize low-latency storage or server cluster traffic in a traffic class to receive more bandwidth and restrict best-effort LAN traffic assigned to a different traffic class.

Although you can configure strict-priority queue scheduling for a priority group, ETS introduces flexibility that allows the bandwidth allocated to each priority group to be dynamically managed according to the amount of LAN, storage, and server traffic in a flow. Unused bandwidth is dynamically allocated to prioritized priority groups. Traffic is queued according to its 802.1p priority assignment, while flexible bandwidth allocation and the configured queue-scheduling for a priority group is supported.

The following figure shows how ETS allows you to allocate bandwidth when different traffic types are classed according to 802.1p priority and mapped to priority groups.




Figure 28. Enhanced Transmission Selection

The following table lists the traffic groupings ETS uses to select multiprotocol traffic for transmission.

Table 9. ETS Traffic Groupings

Traffic Groupings	Description
Priority group	A group of 802.1p priorities used for bandwidth allocation and queue scheduling. All 802.1p priority traffic in a group must have the same traffic handling requirements for latency and frame loss.
Group ID	A 4-bit identifier assigned to each priority group. The range is from 0 to 7.
Group bandwidth	Percentage of available bandwidth allocated to a priority group.
Group transmission selection algorithm (TSA)	Type of queue scheduling a priority group uses.

In the Dell Networking OS, ETS is implemented as follows:

- ETS supports groups of 802.1p priorities that have:
 - PFC enabled or disabled
 - No bandwidth limit or no ETS processing
- Bandwidth allocated by the ETS algorithm is made available after strict-priority groups are serviced. If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups.
- For ETS traffic selection, an algorithm is applied to priority groups using:
 - Strict priority shaping
 - ETS shaping
- ETS uses the DCB MIB IEEE 802.1azd2.5.

Data Center Bridging Exchange Protocol (DCBx)

DCBx allows a switch to automatically discover DCB-enabled peers and exchange configuration information. PFC and ETS use DCBx to exchange and negotiate parameters with peer devices. DCBx capabilities include:

- Discovery of DCB capabilities on peer-device connections.
- Determination of possible mismatch in DCB configuration on a peer link.
- Configuration of a peer device over a DCB link.

DCBx requires the link layer discovery protocol (LLDP) to provide the path to exchange DCB parameters with peer devices. Exchanged parameters are sent in organizationally specific TLVs in LLDP data units. For more information, refer to [Link Layer Discovery Protocol \(LLDP\)](#). The following LLDP TLVs are supported for DCB parameter exchange:

PFC parameters PFC Configuration TLV and Application Priority Configuration TLV.

ETS parameters ETS Configuration TLV and ETS Recommendation TLV.

Data Center Bridging in a Traffic Flow

The following figure shows how DCB handles a traffic flow on an interface.

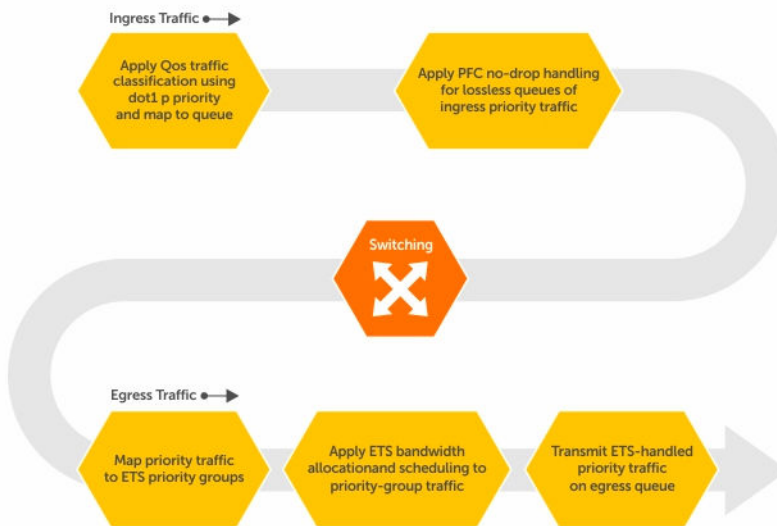


Figure 29. DCB PFC and ETS Traffic Handling

Enabling Data Center Bridging

Data center bridging is enabled by default on an MXL 10/40GbE Switch to support converged enhanced Ethernet (CEE) in a data center network.

A prerequisite for configuring DCB:

- Priority-based flow control
- Enhanced transmission selection
- Data center bridging exchange protocol
- FCoE initialization protocol (FIP) snooping

DCB processes virtual local area network (VLAN)-tagged packets and dot1p priority values. Untagged packets are treated with a dot1p priority of 0.

For DCB to operate effectively, you can classify ingress traffic according to its dot1p priority so that it maps to different data queues. The dot1p-queue assignments used are shown in the following table.


On the MXL Switch, by default, DCB is enabled and MMU buffers are reserved to achieve no-drop traffic handling for PFC. Disabling DCB does not release the buffers reserved by default. To utilize reserved buffers for non-DCB applications, you have to explicitly release the buffers (Refer to [Configuring the PFC Buffer in a Switch Stack](#)).

To disable or re-enable DCB on a switch, enter the following commands.

1. Disable DCB.
CONFIGURATION mode

`no dcb enable`
2. Re-enable DCB.
CONFIGURATION mode

`dcb enable`

 **NOTE: Dell Networking OS Behavior:** DCB is not supported if you enable link-level flow control on one or more interfaces.

After you disable DCB, if link-level flow control is not automatically enabled on an interface, to enable flow control, manually shut down the interface (the `shutdown` command) and re-enable it (the `no shutdown` command).

Configuring DCB Maps and its Attributes

This topic contains the following sections that describe how to configure a DCB map, apply the configured DCB map to a port, configure PFC without a DCB map, and configure lossless queues.

DCB Map: Configuration Procedure

A DCB map consists of PFC and ETS parameters. By default, PFC is not enabled on any 802.1p priority and ETS allocates equal bandwidth to each priority. To configure user-defined PFC and ETS settings, you must create a DCB map.

Step	Task	Command	Command Mode
1	Enter global configuration mode to create a DCB map or edit PFC and ETS settings.	<code>dcb-map name</code>	CONFIGURATION
2	Configure the PFC setting (on or off) and the ETS bandwidth percentage allocated to traffic in each priority group, or whether the priority group traffic should be handled with strict priority scheduling. You can enable PFC on a maximum of two priority queues on an interface. Enabling PFC for dot1p priorities makes the corresponding port queue lossless. The sum of all allocated bandwidth percentages in all groups in the DCB map must be 100%. Strict-priority traffic is serviced first. Afterwards, bandwidth allocated to other priority groups is made available and allocated according to the specified percentages. If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups. Example: priority-group 0 bandwidth 60 pfc off priority-group 1 bandwidth 20 pfc on priority-group 2 bandwidth 20 pfc on priority-group 4 strict-priority pfc off Repeat this step to configure PFC and ETS traffic handling for each priority group.	<code>priority-group group_num {bandwidth percentage strict-priority} pfc {on off}</code>	DCB MAP
3	Specify the dot1p priority-to-priority group mapping for each priority. Priority-group range: 0 to 7. All priorities that map to the same queue must be in the same priority group. Leave a space between each priority group number. For example: priority-pgid 0 0 0 1 2 4 4 4 in which priority group 0 maps to dot1p priorities 0, 1, and 2; priority group 1 maps to dot1p priority 3; priority group 2 maps to dot1p priority 4; priority group 4 maps to dot1p priorities 5, 6, and 7.	<code>priority-pgid dot1p0_group_num dot1p1_group_num dot1p2_group_num dot1p3_group_num dot1p4_group_num dot1p5_group_num dot1p6_group_num dot1p7_group_num</code>	DCB MAP

Important Points to Remember

- If you remove a dot1p priority-to-priority group mapping from a DCB map (`no priority pgid` command), the PFC and ETS parameters revert to their default values on the interfaces on which the DCB map is applied. By default, PFC is not applied on specific 802.1p priorities; ETS assigns equal bandwidth to each 802.1p priority.

As a result, PFC and lossless port queues are disabled on 802.1p priorities, and all priorities are mapped to the same priority queue and equally share the port bandwidth.

- To change the ETS bandwidth allocation configured for a priority group in a DCB map, do not modify the existing DCB map configuration. Instead, first create a new DCB map with the desired PFC and ETS settings, and apply the new map to the interfaces to override the previous DCB map settings. Then, delete the original dot1p priority-priority group mapping.

If you delete the dot1p priority-priority group mapping (`no priority pgid` command) before you apply the new DCB map, the default PFC and ETS parameters are applied on the interfaces. This change may create a DCB mismatch with peer DCB devices and interrupt network operation.

Applying a DCB Map on a Port

To apply a DCB map to an Ethernet port, follow these steps:

Step	Task	Command	Command Mode
1	Enter interface configuration mode on an Ethernet port.	<code>interface</code> <code>{tengigabitEthernet slot/</code> <code>port </code> <code>fortygigabitEthernet</code> <code>slot/port}</code>	CONFIGURATION
2	Apply the DCB map on the Ethernet port to configure it with the PFC and ETS settings in the map; for example: Dell# interface tengigabitEthernet 0/0 Dell(config-if-te-0/0)# dcb-map SAN_A_dcb_map1 Repeat Steps 1 and 2 to apply a DCB map to more than one port. You cannot apply a DCB map on an interface that has been already configured for PFC using the <code>pfc priority</code> command or which is already configured for lossless queues (<code>pfc no-drop queues</code> command).	<code>dcb-map name</code>	INTERFACE

Configuring PFC without a DCB Map

In a network topology that uses the default ETS bandwidth allocation (assigns equal bandwidth to each priority), you can also enable PFC for specific dot1p-priorities on individual interfaces without using a DCB map. This type of DCB configuration is useful on interfaces that require PFC for lossless traffic, but do not transmit converged Ethernet traffic.

Step	Task	Command	Command Mode
1	Enter interface configuration mode on an Ethernet port.	<code>interface {tengigabitEthernet slot/port fortygigabitEthernet slot/port}</code>	CONFIGURATION
2	Enable PFC on specified priorities. Range: 0-7. Default: None. Maximum number of lossless queues supported on an Ethernet port: 2. Separate priority values with a comma. Specify a priority range with a dash, for example: <code>pfc priority 3,5-7</code> 1. You cannot configure PFC using the <code>pfc priority</code> command on an interface on which a DCB map has been applied or which is already configured for lossless queues (<code>pfc no-drop queues</code> command).	<code>pfc priority priority-range</code>	INTERFACE

Configuring Lossless Queues

DCB also supports the manual configuration of lossless queues on an interface after you disable PFC mode in a DCB map and apply the map on the interface. The configuration of no-drop queues provides flexibility for ports on which PFC is not needed, but lossless traffic should egress from the interface.

Lossless traffic egresses out the no-drop queues. Ingress 802.1p traffic from PFC-enabled peers is automatically mapped to the no-drop egress queues.

When configuring lossless queues on a port interface, consider the following points:

- By default, no lossless queues are configured on a port.
- A limit of two lossless queues are supported on a port. If the number of lossless queues configured exceeds the maximum supported limit per port (two), an error message is displayed. You must re-configure the value to a smaller number of queues.
- If you configure lossless queues on an interface that already has a DCB map with PFC enabled (**pfc on**), an error message is displayed.

Step	Task	Command	Command Mode
1	Enter INTERFACE Configuration mode.	<code>interface{tengigabitE thernet slot/port fortygigabitEthernet slot/port}</code>	CONFIGURATION
2	Open a DCB map and enter DCB map configuration mode.	<code>dcb-map name</code>	INTERFACE
3	Disable PFC.	<code>no pfc mode on</code>	DCB MAP

Step	Task	Command	Command Mode
4	Return to interface configuration mode.	<code>exit</code>	DCB MAP
5	Apply the DCB map, created to disable the PFC operation, on the interface	<code>dcb-map {name default}</code>	INTERFACE
6	Configure the port queues that still function as no-drop queues for lossless traffic. The maximum number of lossless queues globally supported on a port is 2. You cannot configure PFC no-drop queues on an interface on which a DCB map with PFC enabled has been applied, or which is already configured for PFC using the <code>pfc priority</code> command. Range: 0-3. Separate queue values with a comma; specify a priority range with a dash; for example: <code>pfc no-drop queues 1,3</code> or <code>pfc no-drop queues 2-3</code> Default: No lossless queues are configured.	<code>pfc no-drop queues queue-range</code>	INTERFACE


Data Center Bridging: Default Configuration

Before you configure PFC and ETS on a switch see the priority group setting taken into account the following default settings:

DCB is enabled.

PFC and ETS are globally enabled by default.

The default dot1p priority-queue assignments are applied as follows:

 **NOTE:** In Egress queue assignment (8 queues in S6000 and Z9500, 4 against in S5000 and S4810.

PFC is not applied on specific dot1p priorities.

ETS: Equal bandwidth is assigned to each port queue and each dot1p priority in a priority group.

To configure PFC and ETS parameters on an S6000 interface, you must specify the PFC mode, the ETS bandwidth allocation for a priority group, and the 802.1p priority-to-priority group mapping in a DCB map. No default PFC and ETS settings are applied to Ethernet interfaces.

Interworking of DCB Map With DCB Buffer Threshold Settings

The `dcb-input` and `dcb-output` configuration commands are deprecated. You must use the `dcb-map` command to create a DCB map to configure priority flow control (PFC) and enhanced transmission selection (ETS) on Ethernet ports that support converged Ethernet traffic.

Configure the `dcb-buffer-threshold` command and its related parameters only on ports with either auto configuration or `dcb-map` configuration. This command is not supported on existing front-panel interfaces or stack ports that are configured with the `dcb-input` or `dcb-output` commands. Similarly, if the `dcb-buffer-threshold` configuration is present on a stack port or any interface, the `dcb-input` or `dcb-output` policies cannot be applied on those interfaces.

Example: When the `dcb-buffer-threshold` policy is applied on interfaces or stack ports with the `dcb-input` or `dcb-output` policies, the following error message is displayed:

```
%Error: dcb-buffer-threshold not supported on interfaces with deprecated
commands
```

Example: When the `dcb-input` or `dcb-output` policy is configured on interfaces or stack ports with the `dcb-buffer` threshold policy, the following error message is displayed:

```
%Error: Deprecated command is not supported on interfaces with dcb-buffer-
threshold configured
```

You must not modify the `service-class dot1p` mappings when any `buffer-threshold-policy` is configured on the system.

```
Dell(conf)#service-class dot1p-mapping dot1p0 3
```

```
% Error: PFC buffer-threshold policies conflict with dot1p mappings. Please
remove all dcb-buffer-threshold policies to change mappings.
```

The `show dcb` command has been enhanced to display the following additional buffer-related information:

```
Dell(conf)#do show dcb
dcb Status : Enabled
PFC Queue Count : 2 --Indicate the PFC queue configured.
Total buffer (lossy + lossless) (in KB): 7787--Total buffer space for lossy and
lossless queues
PFC total buffer (in KB): 6526 --Indicates the total buffer (configured or
default)
PFC shared buffer (in KB): 832--Indicates the shared buffer (Configured or
default)
PFC available buffer ( in KB): 5694--Indicates remaining available buffers for
PFC that are free to be allocated
```


Configuring Priority-Based Flow Control

PFC provides a flow control mechanism based on the 802.1p priorities in converged Ethernet traffic received on an interface and is enabled by default when you enable DCB.

As an enhancement to the existing Ethernet pause mechanism, PFC stops traffic transmission for specified priorities (Class of Service (CoS) values) without impacting other priority classes. Different traffic types are assigned to different priority classes.

When traffic congestion occurs, PFC sends a pause frame to a peer device with the CoS priority values of the traffic that is to be stopped. Data Center Bridging Exchange protocol (DCBx) provides the link-level exchange of PFC parameters between peer devices. PFC allows network administrators to create zero-loss links for Storage Area Network (SAN) traffic that requires no-drop service, while retaining packet-drop congestion management for Local Area Network (LAN) traffic.

To configure PFC, follow these steps:

1. Create a DCB Map.

CONFIGURATION mode

```
dcb-map dcb-map-name
```

The *dcb-map-name* variable can have a maximum of 32 characters.

2. Create a PFC group.

CONFIGURATION mode

```
priority-group group-num {bandwidth bandwidth | strict-priority} pfc on
```

The range for priority group is from 0 to 7.

Set the bandwidth in percentage. The percentage range is from 1 to 100% in units of 1%.

Committed and peak bandwidth is in megabits per second. The range is from 0 to 40000.

Committed and peak burst size is in kilobytes. Default is 50. The range is from 0 to 10000.

The `pfc on` command enables priority-based flow control.

3. Specify the dot1p priority-to-priority group mapping for each priority.

```
priority-pgid dot1p0_group_num dot1p1_group_num ...dot1p7_group_num
```

Priority group range is from 0 to 7. All priorities that map to the same queue must be in the same priority group.

Leave a space between each priority group number. For example: `priority-pgid 0 0 0 1 2 4 4 4` in which priority group 0 maps to dot1p priorities 0, 1, and 2; priority group 1 maps to dot1p priority 3; priority group 2 maps to dot1p priority 4; priority group 4 maps to dot1p priorities 5, 6, and 7.

Dell Networking OS Behavior: As soon as you apply a DCB policy with PFC enabled on an interface, DCBx starts exchanging information with PFC-enabled peers. The IEEE802.1Qbb, CEE, and CIN versions of PFC Type, Length, Value (TLV) are supported. DCBx also validates PFC configurations that are received in TLVs from peer devices.

To honor a PFC pause frame multiplied by the number of PFC-enabled ingress ports, the minimum link delay must be greater than the round-trip transmission time the peer requires.

 **NOTE:** You cannot enable PFC and link-level flow control at the same time on an interface.

The Dell Networking OS does not support MACsec Bypass Capability (MBC).

Configuring Lossless Queues

DCB also supports the manual configuration of lossless queues on an interface when PFC mode is turned off and priority classes are disabled in a DCB map, apply the map on the interface.

Prerequisite: A DCB input policy with PFC configuration is applied to the interface with the following conditions:

- PFC mode is off (no `pfc mode on`).
- No PFC priority classes are configured (no `pfc priority priority-range`).

Example:

```
Port A -> Port B
```

```
Port C -> Port B
```

PFC no-drop queues are configured for queues 1, 2 on Port B. PFC capability is enabled on priorities 3, 4 on PORT A and C.

Port B acting as Egress

During the congestion, [traffic pump on priorities 3 and 4 from PORT A and PORT C is at full line rate], PORT A and C send out the PFCs to rate the traffic limit. Egress drops are not observed on Port B since traffic flow on priorities is mapped to loss less queues.

Port B acting as Ingress

If the traffic congestion is on PORT B , Egress DROP is on PORT A or C, as the PFC is not enabled on PORT B.

Refer the following configuration for queue to dot1p mapping:

```
Dell(conf)#do show qos dot1p-queue-mapping
Dot1p Priority : 0 1 2 3 4 5 6 7 -> On ingress interfaces[Port A and
C] we used the PFC on priority level.
Queue : 0 0 0 1 2 3 3 3 -> On Egress interface[Port B] we
used no-drop queues.
```

The configuration of no-drop queues provides flexibility for ports on which PFC is not needed but lossless traffic should egress from the interface.

Lossless traffic egresses out the no-drop queues. Ingress dot1p traffic from PFC-enabled interfaces is automatically mapped to the no-drop egress queues.

1. Enter INTERFACE Configuration mode.
CONFIGURATION mode

```
interface type slot/port
```

2. Configure the port queues that still functions as no-drop queues for lossless traffic.

INTERFACE mode

```
pfcc no-drop queues queue-range
```

For the dot1p-queue assignments, refer to the dot1p Priority-Queue Assignment table.

The maximum number of lossless queues globally supported on the switch is four.

The range is from 0 to 3. Separate the queue values with a comma; specify a priority range with a dash; for example, `pfcc no-drop queues 1,3` or `pfcc no-drop queues 2-3`.

The default: No lossless queues are configured.

 **NOTE: Dell Networking OS Behavior:** By default, no lossless queues are configured on a port.

A limit of two lossless queues is supported on a port. If the amount of priority traffic that you configure to be paused exceeds the two lossless queues, an error message displays. Reconfigure the input policy using a smaller number of PFC priorities.

If you configure lossless queues on an interface that already has a DCB input policy with PFC enabled (`pfcc mode on`), an error message displays.

Traffic may be interrupted due to an interface flap (going down and coming up) when you reconfigure lossless queues on no-drop priorities in an input policy and re-apply the policy to an interface.

Configuring the PFC Buffer in a Switch Stack

In a switch stack, you must configure all stacked ports with the same PFC configuration. In addition, you must configure a separate buffer of memory allocated exclusively to a service pool accessed by queues on which priority-based control flows are mapped.

These PFC-enabled queues ensure the lossless transmission of storage and server traffic. The buffer required for the PFC service pool is calculated based on the number of ports and port queues used by PFC traffic.

You can configure the size of the PFC buffer for all switches in a stack or all port pipes on a specified stack unit by entering the following commands on the master switch.

- Configure the PFC buffer for all switches in the stack.

CONFIGURATION mode

```
[no] dcb stack-unit all pfcc-buffering pfcc-port {1-56} pfcc-queues {1-2}
```

By default, the PFC buffer is enabled on all ports on the stack unit.

- Configure the PFC buffer for all port pipes in a specified stack unit by specifying the port-pipe number, number of PFC-enabled ports, and number of configured lossless queues.

CONFIGURATION mode

```
[no] dcb stack-unit stack-unit-id [port-set port-set-id] pfcc-buffering pfcc-ports {1-56} pfcc-queues {1-2}
```

Valid stack-unit IDs are 0 to 5.

The only valid port-set ID (port-pipe number) is 0.

Dell Networking OS Behavior: If you configure PFC on a 40GbE port, count the 40GbE port as four PFC-enabled ports in the `pfc-port` number you enter in the command syntax.

To achieve lossless PFC operation, the PFC port count and queue number used for the reserved buffer size that is created must be greater than or equal to the buffer size required for PFC-enabled ports and lossless queues on the switch.

For the PFC buffer configuration to take effect, you must reload the stack or a specified stack unit (use the `reload` command at EXEC Privilege level).

If you configure the PFC buffer on all stack units, delete the startup configuration on both the master and standby, and reload the stack, the new master (previously standby) generates the following syslog message for each stack unit when it boots up: `PFC_BUFFER_CONFIG_CHANGED` is generated for all stack units.

Priority-Based Flow Control Using Dynamic Buffer Method

In a data center network, priority-based flow control (PFC) manages large bursts of one traffic type in multiprotocol links so that it does not affect other traffic types and no frames are lost due to congestion. When PFC detects congestion on a queue for a specified priority, it sends a pause frame for the 802.1p priority traffic to the transmitting device.

Pause and Resume of Traffic

The pause message is used by the sending device to inform the receiving device about a congested, heavily-loaded traffic state that has been identified. When the interface of a sending device transmits a pause frame, the recipient acknowledges this frame by temporarily halting the transmission of data packets. The sending device requests the recipient to restart the transmission of data traffic when the congestion eases and reduces. The time period that is specified in the pause frame defines the duration for which the flow of data packets is halted. When the time period elapses, the transmission restarts.

When a device sends a pause frame to another device, the time for which the sending of packets from the other device must be stopped is contained in the pause frame. The device that sent the pause frame empties the buffer to be less than the threshold value and restarts the acceptance of data packets.

Dynamic ingress buffering enables the sending of pause frames at different thresholds based on the number of ports that experience congestion at a time. This behavior impacts the total buffer size used by a particular lossless priority on an interface. The pause and resume thresholds can also be configured dynamically. You can configure a buffer size, pause threshold, ingress shared threshold weight, and resume threshold to control and manage the total amount of buffers that are to be used in your network environment.

All the PFC-related settings such as the DCB input and output policies or DCB maps are saved in the DCB application and the Differentiated Services Manager (DSM) application. All of these configurations can be modified only for interfaces that are enabled for DCB. The DCB buffer configurations are also saved in the DCB and DSM databases.

Buffer Sizes for Lossless or PFC Packets

You can configure up to a maximum of 4 lossless (PFC) queues. By configuring 4 lossless queues, you can configure 4 different priorities and assign a particular priority to each application that your network is

used to process. For example, you can assign a higher priority for time-sensitive applications and a lower priority for other services, such as file transfers. You can configure the amount of buffer space to be allocated for each priority and the pause or resume thresholds for the buffer. This method of configuration enables you to effectively manage and administer the behavior of lossless queues.

Although the system contains 9 MB of space for shared buffers, a minimum guaranteed buffer is provided to all the internal and external ports in the system for both unicast and multicast traffic. This minimum guaranteed buffer reduces the total available shared buffer to 7,787 KB. This shared buffer can be used for lossy and lossless traffic.

The default behavior causes up to a maximum of 6.6 MB to be used for PFC-related traffic. The remaining approximate space of 1 MB can be used by lossy traffic. You can allocate all the remaining 1 MB to lossless PFC queues. If you allocate in such a way, the performance of lossy traffic is reduced and degraded. Although you can allocate a maximum buffer size, it is used only if a PFC priority is configured and applied on the interface.

The number of lossless queues supported on the system is dependent on the availability of total buffers for PFC. The default configuration in the system guarantees a minimum of 52 KB per queue if all the 128 queues are congested. However, modifying the buffer allocation per queue impacts this default behavior.

By default the total available buffer for PFC is 6.6 MB and when you configure dynamic ingress buffering, a minimum of least 52 KB per queue is used when all ports are congested. By default, the system enables a maximum of two lossless queues on the MXL platform.

This default behavior is impacted if you modify the total buffer available for PFC or assign static buffer configurations to the individual PFC queues.

Configure Enhanced Transmission Selection

ETS provides a way to optimize bandwidth allocation to outbound 802.1p classes of converged Ethernet traffic.

Different traffic types have different service needs. Using ETS, you can create groups within an 802.1p priority class to configure different treatment for traffic with different bandwidth, latency, and best-effort needs.

For example, storage traffic is sensitive to frame loss; interprocess communication (IPC) traffic is latency-sensitive. ETS allows different traffic types to coexist without interruption in the same converged link by:

- Allocating a guaranteed share of bandwidth to each priority group.
- Allowing each group to exceed its minimum guaranteed bandwidth if another group is not fully using its allotted bandwidth.


To configure ETS and apply an ETS output policy to an interface, you must:

1. Create a Quality of Service (QoS) output policy with ETS scheduling and bandwidth allocation settings.
2. Create a priority group of 802.1p traffic classes.
3. Configure a DCB output policy in which you associate a priority group with a QoS ETS output policy.
4. Apply the DCB output policy to an interface.

ETS Prerequisites and Restrictions

The following prerequisites and restrictions apply when you configure ETS bandwidth allocation or queue scheduling and apply a QoS ETS output policy on an interface.

- Configuring ETS bandwidth allocation or a queue scheduler for dot1p priorities in a priority group is applicable if the DCBx version used on a port is CIN (refer to [Configuring DCBx](#)) or CEE as a port version where CNA supports CEE and DUT port versions in AUTO or CEE mode.
- When allocating bandwidth or configuring a queue scheduler for dot1p priorities in a priority group on a DCBx CIN interface, take into account the CIN bandwidth allocation (refer to [Configuring Bandwidth Allocation for DCBx CIN](#)) and dot1p-queue mapping ([QoS dot1p Traffic Classification and Queue Assignment](#)).

 **NOTE:** The IEEE 802.1Qaz, CEE, and CIN versions of ETS are supported.

Creating an ETS Priority Group

An ETS priority group specifies the range of 802.1p priority traffic to which a QoS output policy with ETS settings is applied on an egress interface.

1. Configure a DCB Map.

CONFIGURATION mode

```
dcb-map dcb-map-name
```

The *dcb-map-name* variable can have a maximum of 32 characters.

2. Create an ETS priority group.

CONFIGURATION mode

```
priority-group group-num {bandwidth bandwidth | strict-priority} pfc off
```

The range for priority group is from 0 to 7.

Set the bandwidth in percentage. The percentage range is from 1 to 100% in units of 1%.

Committed and peak bandwidth is in megabits per second. The range is from 0 to 40000.

Committed and peak burst size is in kilobytes. Default is 50. The range is from 0 to 10000.

3. Configure the 802.1p priorities for the traffic on which you want to apply an ETS output policy.

PRIORITY-GROUP mode

```
priority-list value
```

The range is from 0 to 7.

The default is **none**.

Separate priority values with a comma. Specify a priority range with a dash. For example, priority-list 3,5-7.

4. Exit priority-group configuration mode.

PRIORITY-GROUP mode

exit

5. Repeat Steps 1 to 4 to configure all remaining dot1p priorities in an ETS priority group.

Dell Networking OS Behavior: A priority group consists of 802.1p priority values that are grouped for similar bandwidth allocation and scheduling, and that share latency and loss requirements. All 802.1p priorities mapped to the same queue must be in the same priority group.

Configure all 802.1p priorities in priority groups associated with an ETS output policy. You can assign each dot1p priority to only one priority group.

By default, all 802.1p priorities are grouped in priority group 0 and 100% of the port bandwidth is assigned to priority group 0. The complete bandwidth is equally assigned to each priority class so that each class has 12 to 13%.

The maximum number of priority groups supported in ETS output policies on an interface is equal to the number of data queues (4) on the port. The 802.1p priorities in a priority group can map to multiple queues.

If you configure more than one priority queue as strict priority or more than one priority group as strict priority, the higher numbered priority queue is given preference when scheduling data traffic.

ETS Operation with DCBx

The following section describes DCBx negotiation with peer ETS devices.

In DCBx negotiation with peer ETS devices, ETS configuration is handled as follows:

- ETS TLVs are supported in DCBx versions CIN, CEE, and IEEE2.5.
- The DCBx port-role configurations determine the ETS operational parameters (refer to [Configure a DCBx Operation](#)).
- ETS configurations received from TLVs from a peer are validated.
- If there is a hardware limitation or TLV error:
 - DCBx operation on an ETS port goes down.
 - New ETS configurations are ignored and existing ETS configurations are reset to the previously configured ETS output policy on the port or to the default ETS settings if no ETS output policy was previously applied.
- ETS operates with legacy DCBx versions as follows:
 - In the CEE version, the priority group/traffic class group (TCG) ID 15 represents a non-ETS priority group. Any priority group configured with a scheduler type is treated as a strict-priority group and is given the priority-group (TCG) ID 15.
 - The CIN version supports two types of strict-priority scheduling:
 - * Group strict priority: Use this to increase its bandwidth usage to the bandwidth total of the priority group and allow a single priority flow in a priority group. A single flow in a group can use all the bandwidth allocated to the group.
 - * Link strict priority: Use this to increase to the maximum link bandwidth and allow a flow in any priority group.

CIN supports only the dot1p priority-queue assignment in a priority group. To configure a dot1p priority flow in a priority group to operate with link strict priority, you configure: The dot1p priority for strict-priority scheduling (`strict-priority` command; [Enabling Strict-Priority Queueing](#)).

If you configure only the priority group in an ETS output policy or only the dot1p priority for strict-priority scheduling, the flow is handled with group strict priority.

Configuring Bandwidth Allocation for DCBx CIN

After you apply an ETS output policy to an interface, if the DCBx version used in your data center network is CIN, you may need to configure a QoS output policy to overwrite the default CIN bandwidth allocation. This default setting divides the bandwidth allocated to each port queue equally between the dot1p priority traffic assigned to the queue.

For more information, refer to [Allocating Bandwidth to Queue](#).

To create a QoS output policy that allocates different amounts of bandwidth to the different traffic types/dot1p priorities assigned to a queue and apply the output policy to the interface, follow these steps.

1. Create a QoS output policy.

CONFIGURATION mode

```
Dell(conf)#qos-policy-output test12
```

The maximum 32 alphanumeric characters.

2. Configure the percentage of bandwidth to allocate to the dot1p priority/queue traffic in the associated L2 class map.

QoS OUTPUT POLICY mode

```
Dell(conf-qos-policy-out)#bandwidth-percentage 100
```

The default is **none**.

3. Repeat Step 2 to configure bandwidth percentages for other priority queues on the port.

QoS OUTPUT POLICY mode

```
Dell(conf-qos-policy-out)#bandwidth-percentage 100
```

4. Create a priority group for strict-priority scheduling.

QoS OUTPUT POLICY mode

```
Dell(conf-qos-policy-out)#scheduler strict
```



NOTE: You can not use `scheduler strict` when bandwidth percentage is configured. It displays an error message.

```
Dell(conf-qos-policy-out)#bandwidth-percentage 100
Dell(conf-qos-policy-out)#scheduler strict
% Error: Strict priority scheduler mode is not allowed when bandwidth-
percentage is configured on qos-policy-output profile.
Dell(conf-qos-policy-out)#scheduler strict ?
```

5. Exit QoS Output Policy Configuration mode.

QoS OUTPUT POLICY mode

```
Dell(conf-if-te-0/1)#exit
```

6. Enter INTERFACE Configuration mode.

CONFIGURATION mode

```
Dell(conf-qos-policy-out)#int te 0/1
```


7. Apply the QoS output policy with the bandwidth percentage for specified priority queues to an egress interface.

INTERFACE mode

```
Dell(conf-if-te-0/1)#service-policy output test12
```

Hierarchical Scheduling in ETS Output Policies

ETS supports up to three levels of hierarchical scheduling.

For example, you can apply ETS output policies with the following configurations:

- | | |
|-------------------------|---|
| Priority group 1 | Assigns traffic to one priority queue with 20% of the link bandwidth and strict-priority scheduling. |
| Priority group 2 | Assigns traffic to one priority queue with 30% of the link bandwidth. |
| Priority group 3 | Assigns traffic to two priority queues with 50% of the link bandwidth and strict-priority scheduling. |

In this example, the configured ETS bandwidth allocation and scheduler behavior is as follows:

- | | |
|--------------------------------|---|
| Unused bandwidth usage: | <p>Normally, if there is no traffic or unused bandwidth for a priority group, the bandwidth allocated to the group is distributed to the other priority groups according to the bandwidth percentage allocated to each group. However, when three priority groups with different bandwidth allocations are used on an interface:</p> <ul style="list-style-type: none">• If priority group 3 has free bandwidth, it is distributed as follows: 20% of the free bandwidth to priority group 1 and 30% of the free bandwidth to priority group 2.• If priority group 1 or 2 has free bandwidth, (20 + 30)% of the free bandwidth is distributed to priority group 3. Priority groups 1 and 2 retain whatever free bandwidth remains up to the (20+ 30)%. |
| Strict-priority groups: | <p>If two priority groups have strict-priority scheduling, traffic assigned from the priority group with the higher priority-queue number is scheduled first. However, when three priority groups are used and two groups have strict-priority scheduling (such as groups 1 and 3 in the example), the strict priority group whose traffic is mapped to one queue takes precedence over the strict priority group whose traffic is mapped to two queues.</p> |

Therefore, in this example, scheduling traffic to priority group 1 (mapped to one strict-priority queue) takes precedence over scheduling traffic to priority group 3 (mapped to two strict-priority queues).

Applying DCB Policies with an ETS Configuration

You can apply a DCB output policy with ETS configuration to all stacked ports in a switch stack or an individual stacked switch. In addition, you can apply different DCB output policies to different stack units.

- Apply the specified DCB output policy on all ports of the switch stack or a stacked switch.
CONFIGURATION mode

```
dcb-policy output stack-unit {all | stack-unit-id} stack-ports all dcb-  
output-policy-name
```

Entering this command removes all DCB input policies applied to stacked ports.

Dell Networking Behavior: A `dcb-policy output stack-unit all` command overwrites any previous `dcb-policy output stack-unit stack-unit-id` configurations. Similarly, a `dcb-policy output stack-unit stack-unit-id` command overwrites any previous `dcb-policy output stack-unit all` configuration.

Entering the `no dcb-policy output stack-unit all` command removes all DCB output policies applied to stacked ports. The `no dcb-policy output stack-unit stack-unit-id` command removes only the DCB output policy applied to the specified switch.

PFC and ETS Configuration Examples

This section contains examples of how to configure and apply DCB input and output policies on an interface.

Using PFC and ETS to Manage Data Center Traffic

The following shows examples of using PFC and ETS to manage your data center traffic.

In the following example:

- Incoming SAN traffic is configured for priority-based flow control.
- Outbound LAN, IPC, and SAN traffic is mapped into three ETS priority groups and configured for enhanced traffic selection (bandwidth allocation and scheduling).
- One lossless queue is used.




Figure 30. PFC and ETS Applied to LAN, IPC, and SAN Priority Traffic

QoS Traffic Classification: The `service-class dynamic dot1p` command has been used in Global Configuration mode to map ingress dot1p frames to the queues shown in the following table. For more information, refer to [QoS dot1p Traffic Classification and Queue Assignment](#).

dot1p Value in Incoming Frame	Queue Assignment
0	0
1	0
2	0
3	1

dot1p Value in Incoming Frame	Queue Assignment
4	2
5	3
6	3
7	3

The following describes the dot1p-priority class group assignment

dot1p Value in the Incoming Frame	Priority Group Assignment
0	LAN
1	LAN
2	LAN
3	SAN
4	IPC
5	LAN
6	LAN
7	LAN

The following describes the priority group-bandwidth assignment.

Priority Group	Bandwidth Assignment
IPC	5%
SAN	50%
LAN	45%

PFC and ETS Configuration Command Examples

The following examples show PFC and ETS configuration commands to manage your data center traffic.

Example of Configuring QoS Priority-Queue Assignment to Honor Dot1p Priorities

```
Dell(conf)# service-class dynamic dot1p
```

Using PFC and ETS to Manage Converged Ethernet Traffic in a Switch Stack

The following example shows how to apply the DCB PFC input policy (`ipc_san_lan`) and ETS output policy (`ets`) on all MXL switches in a switch stack.

This example references the [PFC and ETS Configuration Examples](#) section.

Example of Applying DCB PFC Input Policy and ETS Output Policy in a Switch Stack

```
dcb-map stack-unit all stack-ports all <dcb-map-name>
```

Applying DCB Policies in a Switch Stack

You can apply a DCB policy with PFC configuration to all stacked ports in a switch stack or on a stacked switch. You can apply different DCB policies to different stacked switches. To apply DCB policies in a switch stack, use the following command.

- Apply the specified DCB policy on all ports of the switch stack or a single stacked switch.
CONFIGURATION mode

```
dcb-map stack-unit all stack-ports all <dcb-map-name>
```

Configure a DCBx Operation

DCB devices use data center bridging exchange protocol (DCBx) to exchange configuration information with directly connected peers using the link layer discovery protocol (LLDP) protocol.

DCBx can detect the misconfiguration of a peer DCB device, and optionally, configure peer DCB devices with DCB feature settings to ensure consistent operation in a data center network.

DCBx is a prerequisite for using DCB features, such as priority-based flow control (PFC) and enhanced traffic selection (ETS), to exchange link-level configurations in a converged Ethernet environment. DCBx is also deployed in topologies that support lossless operation for FCoE or iSCSI traffic. In these scenarios, all network devices are DCBx-enabled (DCBx is enabled end-to-end). For more information about how these features are implemented and used, refer to:

- [Configuring Priority-Based Flow Control](#)
- [Configure Enhanced Transmission Selection](#)
- [Configuring FIP Snooping](#)

DCBx supports the following versions: CIN, CEE, and IEEE2.5.

Prerequisite: For DCBx, enable LLDP on all DCB devices.

DCBx Operation

DCBx performs the following operations:

- Discovers DCB configuration (such as PFC and ETS) in a peer device.
- Detects DCB mis-configuration in a peer device; that is, when DCB features are not compatibly configured on a peer device and the local switch. Mis-configuration detection is feature-specific because some DCB features support asymmetric configuration.
- Reconfigures a peer device with the DCB configuration from its configuration source if the peer device is willing to accept configuration.
- Accepts the DCB configuration from a peer if a DCBx port is in “willing” mode to accept a peer’s DCB settings and then internally propagates the received DCB configuration to its peer ports.

DCBx Port Roles

To enable the auto-configuration of DCBx-enabled ports and propagate DCB configurations learned from peer DCBx devices internally to other switch ports, use the following DCBx port roles.

Auto-upstream

The port advertises its own configuration to DCBx peers and receives its configuration from DCBX peers (ToR or FCF device). The port also propagates its configuration to other ports on the switch.

The first auto-upstream that is capable of receiving a peer configuration is elected as the *configuration source*. The elected configuration source then internally propagates the configuration to other auto-upstream and auto-downstream ports. A port that receives an internally propagated configuration overwrites its local configuration with the new parameter values.

When an auto-upstream port (besides the configuration source) receives and overwrites its configuration with internally propagated information, one of the following actions is taken:

- If the peer configuration received is compatible with the internally propagated port configuration, the link with the DCBx peer is enabled.
- If the received peer configuration is not compatible with the currently configured port configuration, the link with the DCBX peer port is disabled and a syslog message for an incompatible configuration is generated. The network administrator must then reconfigure the peer device so that it advertises a compatible DCB configuration.

The configuration received from a DCBX peer or from an internally propagated configuration is not stored in the switch's running configuration.

On a DCBX port in an auto-upstream role, the PFC and application priority TLVs are enabled. ETS recommend TLVs are disabled and ETS configuration TLVs are enabled.

Auto-downstream

The port advertises its own configuration to DCBx peers but is *not willing* to receive remote peer configuration. The port always accepts internally propagated configurations from a configuration source. An auto-downstream port that receives an internally propagated configuration overwrites its local configuration with the new parameter values.

When an auto-downstream port receives and overwrites its configuration with internally propagated information, one of the following actions is taken:

- If the peer configuration received is compatible with the internally propagated port configuration, the link with the DCBx peer is enabled.
- If the received peer configuration is not compatible with the currently configured port configuration, the link with the DCBX peer port is disabled and a syslog message for an incompatible configuration is generated. The network administrator must then reconfigure the peer device so that it advertises a compatible DCB configuration.

The internally propagated configuration is not stored in the switch's running configuration. On a DCBX port in an auto-downstream role, all PFC, application priority, ETS recommend, and ETS configuration TLVs are enabled.

Configuration source

The port is configured to serve as a source of configuration information on the switch. Peer DCB configurations received on the port are propagated to other DCBX auto-configured ports. If the peer configuration is compatible with a port configuration, DCBX is enabled on the port.

On a configuration-source port, the link with a DCBX peer is enabled when the port receives a DCB configuration that can be internally propagated to other auto-configured ports.

The configuration received from a DCBX peer is not stored in the switch's running configuration.

On a DCBX port that is the configuration source, all PFC and application priority TLVs are enabled. ETS recommend TLVs are disabled and ETS configuration TLVs are enabled.

Manual

The port is configured to operate only with administrator-configured settings and does not auto-configure with DCB settings received from a DCBX peer or from an internally propagated configuration from the configuration source. If you enable DCBX, ports in Manual mode advertise their configurations to peer devices but do not accept or propagate internal or external configurations. Unlike other user-configured ports, the configuration of DCBX ports in Manual mode is saved in the running configuration.

On a DCBX port in a manual role, all PFC, application priority, ETS recommend, and ETS configuration TLVs are enabled.

The default for the DCBX port role is **manual**.



NOTE: On a DCBX port, application priority TLV advertisements are handled as follows:

- The application priority TLV is transmitted only if the priorities in the advertisement match the configured PFC priorities on the port.
- On auto-upstream and auto-downstream ports:
 - If a configuration source is elected, the ports send an application priority TLV based on the application priority TLV received on the configuration-source port. When an application priority TLV is received on the configuration-source port, the auto-upstream and auto-downstream ports use the internally propagated PFC priorities to match against the received application priority. Otherwise, these ports use their locally configured PFC priorities in application priority TLVs.
 - If no configuration source is configured, auto-upstream and auto-downstream ports check to see that the locally configured PFC priorities match the priorities in a received application priority TLV.
- On manual ports, an application priority TLV is advertised only if the priorities in the TLV match the PFC priorities configured on the port.

DCB Configuration Exchange

The DCBx protocol supports the exchange and propagation of configuration information for the enhanced transmission selection (ETS) and priority-based flow control (PFC) DCB features.

DCBx uses the following methods to exchange DCB configuration parameters:

- | | |
|-------------------|---|
| Asymmetric | DCB parameters are exchanged between a DCBx-enabled port and a peer port without requiring that a peer port and the local port use the same configured values for the configurations to be compatible. For example, ETS uses an asymmetric exchange of parameters between DCBx peers. |
| Symmetric | DCB parameters are exchanged between a DCBx-enabled port and a peer port but requires that each configured parameter value be the same for the configurations in order to be compatible. For example, PFC uses an symmetric exchange of parameters between DCBx peers. |

Configuration Source Election

When an auto-upstream or auto-downstream port receives a DCB configuration from a peer, the port first checks to see if there is an active configuration source on the switch.

- If a configuration source already exists, the received peer configuration is checked against the local port configuration. If the received configuration is compatible, the DCBx marks the port as DCBx-enabled. If the configuration received from the peer is not compatible, a warning message is logged and the DCBx frame error counter is incremented. Although DCBx is operationally disabled, the port keeps the peer link up and continues to exchange DCBx packets. If a compatible peer configuration is later received, DCBx is enabled on the port.
- If there is no configuration source, a port may elect itself as the configuration source. A port may become the configuration source if the following conditions exist:
 - No other port is the configuration source.
 - The port role is auto-upstream.
 - The port is enabled with link up and DCBx enabled.
 - The port has performed a DCBx exchange with a DCBx peer.
 - The switch is capable of supporting the received DCB configuration values through either a symmetric or asymmetric parameter exchange.


A newly elected configuration source propagates configuration changes received from a peer to the other auto-configuration ports. Ports receiving auto-configuration information from the configuration source ignore their current settings and use the configuration source information.

Propagation of DCB Information

When an auto-upstream or auto-downstream port receives a DCB configuration from a peer, the port acts as a DCBx client and checks if a DCBx configuration source exists on the switch.

- If a configuration source is found, the received configuration is checked against the currently configured values that are internally propagated by the configuration source. If the local configuration is compatible with the received configuration, the port is enabled for DCBx operation and synchronization.

- If the configuration received from the peer is not compatible with the internally propagated configuration used by the configuration source, the port is disabled as a client for DCBx operation and synchronization and a syslog error message is generated. The port keeps the peer link up and continues to exchange DCBx packets. If a compatible configuration is later received from the peer, the port is enabled for DCBx operation.

 **NOTE:** DCB configurations internally propagated from a configuration source do not overwrite the configuration on a DCBx port in a manual role. When a configuration source is elected, all auto-upstream ports other than the configuration source are marked as *willing disabled*. The internally propagated DCB configuration is refreshed on all auto-configuration ports and each port may begin configuration negotiation with a DCBx peer again.


Auto-Detection and Manual Configuration of the DCBx Version

When operating in Auto-Detection mode (the `DCBx version auto` command), a DCBx port automatically detects the DCBx version on a peer port. Legacy CIN and CEE versions are supported in addition to the standard IEEE version 2.5 DCBx.

A DCBx port detects a peer version after receiving a valid frame for that version. The local DCBx port reconfigures to operate with the peer version and maintains the peer version on the link until one of the following conditions occurs:

- The switch reboots.
- The link is reset (goes down and up).
- User-configured CLI commands require the version negotiation to restart.
- The peer times out.
- Multiple peers are detected on the link.

If you configure a DCBx port to operate with a specific version (the `DCBx version {cee | cin | ieee-v2.5}` command in the [Configuring DCBx](#)), DCBx operations are performed according to the configured version, including fast and slow transmit timers and message formats. If a DCBx frame with a different version is received, a syslog message is generated and the peer version is recorded in the peer status table. If the frame cannot be processed, it is discarded and the discard counter is incremented.

 **NOTE:** Because DCBx TLV processing is best effort, it is possible that CIN frames may be processed when DCBx is configured to operate in CEE mode and vice versa. In this case, the unrecognized TLVs cause the unrecognized TLV counter to increment, but the frame is processed and is not discarded.

Legacy DCBx (CIN and CEE) supports the DCBx control state machine that is defined to maintain the sequence number and acknowledge the number sent in the DCBx control TLVs.

DCBx Example

The following figure shows how DCBx is used on an MXL Switch installed in a PowerEdge M1000e chassis in which servers are also installed.

The external 40GbE ports on the base module (ports 33 and 37) of two switches are used for uplinks configured as DCBx auto-upstream ports. The MXL switch is connected to third-party, top-of-rack (ToR) switches through 40GbE uplinks. The ToR switches are part of a Fibre Channel storage network.

The internal ports (ports 1-32) connected to the 10GbE backplane are configured as auto-downstream ports.

On the MXL switch, PFC and ETS use DCBx to exchange link-level configuration with DCBx peer devices.

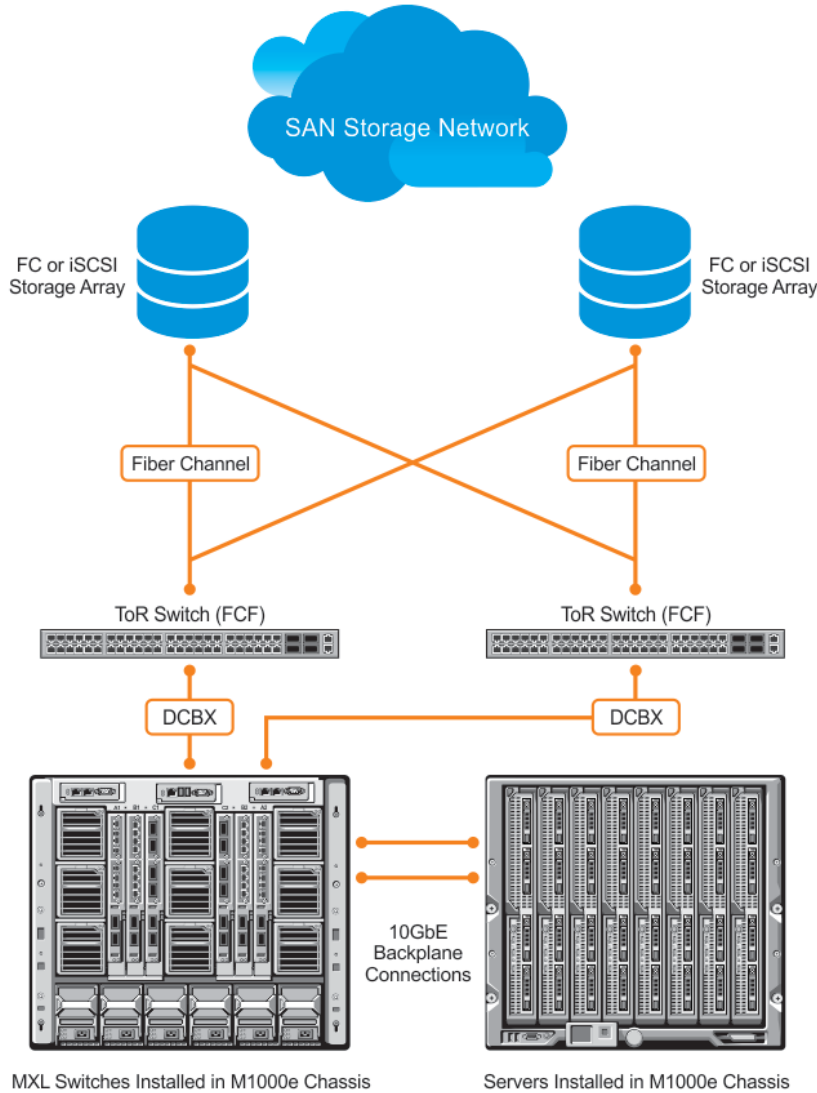


Figure 31. DCBx Sample Topology

DCBx Prerequisites and Restrictions

The following prerequisites and restrictions apply when you configure DCBx operation on a port:

- For DCBx, on a port interface, enable LLDP in both Send (TX) and Receive (RX) mode (the `protocol lldp mode` command; refer to the example in [CONFIGURATION versus INTERFACE Configurations](#) in the [Link Layer Discovery Protocol \(LLDP\)](#) chapter). If multiple DCBx peer ports are detected on a local DCBx interface, LLDP is shut down.
- The CIN version of DCBx supports only PFC, ETS, and FCOE; it does not support iSCSI, backward congestion management (BCN), logical link down (LLDF), and network interface virtualization (NIV).

Configuring DCBx

To configure DCBx, follow these steps.

For DCBx, to advertise DCBx TLVs to peers, enable LLDP. For more information, refer to [Link Layer Discovery Protocol \(LLDP\)](#).

Configure DCBx operation at the interface level on a switch or globally on the switch. To configure an MXL switch for DCBx operation in a data center network, you must:

1. Configure ToR- and FCF-facing interfaces as auto-upstream ports.
2. Configure server-facing interfaces as auto-downstream ports.
3. Configure a port to operate in a configuration-source role.
4. Configure ports to operate in a manual role.

1. Enter INTERFACE Configuration mode.
CONFIGURATION mode

```
interface type slot/port
```

2. Enter LLDP Configuration mode to enable DCBx operation.
INTERFACE mode

```
[no] protocol lldp
```

3. Configure the DCBx version used on the interface, where: `auto` configures the port to operate using the DCBx version received from a peer.

```
PROTOCOL LLDP mode
```

```
[no] DCBx version {auto | cee | cin | ieee-v2.5}
```

- `cee`: configures the port to use CEE (Intel 1.01).
- `cin`: configures the port to use Cisco-Intel-Nuova (DCBx 1.0).
- `ieee-v2.5`: configures the port to use IEEE 802.1Qaz (Draft 2.5).

The default is **Auto**.

4. Configure the DCBx port role the interface uses to exchange DCB information.

```
PROTOCOL LLDP mode
```

```
[no] DCBx port-role {config-source | auto-downstream | auto-upstream | manual}
```

- `auto-upstream`: configures the port to receive a peer configuration. The configuration source is elected from auto-upstream ports.
- `auto-downstream`: configures the port to accept the internally propagated DCB configuration from a configuration source.
- `config-source`: configures the port to serve as the configuration source on the switch.
- `manual`: configures the port to operate only on administer-configured DCB parameters. The port does not accept a DCB configuration received from a peer or a local configuration source.

The default is **Manual**.


5. **On manual ports only:** Configure the PFC and ETS TLVs advertised to DCBx peers.

```
PROTOCOL LLDP mode
```

```
[no] advertise DCBx-tlv {ets-conf | ets-reco | pfc} [ets-conf | ets-reco | pfc] [ets-conf | ets-reco | pfc]
```

- `ets-conf`: enables the advertisement of ETS Configuration TLVs.
- `ets-reco`: enables the advertisement of ETS Recommend TLVs.
- `pfc enables`: the advertisement of PFC TLVs.

The default is All PFC and ETS TLVs are advertised.

 **NOTE:** You can configure the transmission of more than one TLV type at a time; for example, `advertise DCBx-tlv ets-conf ets-reco`. You can enable ETS recommend TLVs (`ets-reco`) only if you enable ETS configuration TLVs (`ets-conf`).

To disable TLV transmission, use the `no` form of the command; for example, `no advertise DCBx-tlv pfc ets-reco`.


6. **On manual ports only:** Configure the Application Priority TLVs advertised on the interface to DCBx peers.

PROTOCOL LLDL mode

```
[no] advertise DCBx-appln-tlv {fcoe | iscsi}
```

- `fcoe`: enables the advertisement of FCoE in Application Priority TLVs.
- `iscsi`: enables the advertisement of iSCSI in Application Priority TLVs.

The default is Application Priority TLVs are enabled to advertise FCoE and iSCSI.

 **NOTE:** To disable TLV transmission, use the `no` form of the command; for example, `no advertise DCBx-appln-tlv iscsi`.

For information about how to use FCoE and iSCSI, refer to [Fibre Channel over Ethernet](#) and [iSCSI Optimization](#).

To verify the DCBx configuration on a port, use the `show interface DCBx detail` command.

Configuring DCBx Globally on the Switch

To globally configure the DCBx operation on a switch, follow these steps.

1. Enter Global Configuration mode.

```
EXEC PRIVILEGE mode
```

```
configure
```

2. Enter LLDP Configuration mode to enable DCBx operation.

```
CONFIGURATION mode
```

```
[no] protocol lldp
```


3. Configure the DCBx version used on all interfaces not already configured to exchange DCB information.

```
PROTOCL LLDL mode
```

```
[no] DCBx version {auto | cee | cin | ieee-v2.5}
```

- `auto`: configures all ports to operate using the DCBx version received from a peer.
- `cee`: configures a port to use CEE (Intel 1.01). `cin` configures a port to use Cisco-Intel-Nuova (DCBx 1.0).
- `ieee-v2.5`: configures a port to use IEEE 802.1Qaz (Draft 2.5).


The default is **Auto**.

 **NOTE:** To configure the DCBx port role the interfaces use to exchange DCB information, use the `DCBx port-role` command in INTERFACE Configuration mode (Step 3).

4. Configure the PFC and ETS TLVs that advertise on unconfigured interfaces with a manual port-role.
PROTOCOL LLDP mode

```
[no] advertise DCBx-tlv {ets-conf | ets-reco | pfc} [ets-conf | ets-reco | pfc] [ets-conf | ets-reco | pfc]
```

- `ets-conf`: enables transmission of ETS Configuration TLVs.
- `ets-reco`: enables transmission of ETS Recommend TLVs.
- `pfc`: enables transmission of PFC TLVs.

 **NOTE:** You can configure the transmission of more than one TLV type at a time. You can only enable ETS recommend TLVs (`ets-reco`) if you enable ETS configuration TLVs (`ets-conf`). To disable TLV transmission, use the `no` form of the command; for example, `no advertise DCBx-tlv pfc ets-reco`.


The default is All TLV types are enabled.

5. Configure the Application Priority TLVs that advertise on unconfigured interfaces with a manual port-role.
PROTOCOL LLDP mode

```
[no] advertise DCBx-appln-tlv {fcoe | iscsi}
```

- `fcoe`: enables the advertisement of FCoE in Application Priority TLVs.
- `iscsi`: enables the advertisement of iSCSI in Application Priority TLVs.

The default is Application Priority TLVs are enabled and advertise FCoE and iSCSI.

 **NOTE:** To disable TLV transmission, use the `no` form of the command; for example, `no advertise DCBx-appln-tlv iscsi`.

For information about how to use FCoE and iSCSI, refer to [Fibre Channel over Ethernet](#) and [iSCSI Optimization](#).

6. Configure the FCoE priority advertised for the FCoE protocol in Application Priority TLVs.
PROTOCOL LLDP mode

```
[no] fcoe priority-bits priority-bitmap
```

The `priority-bitmap` range is from 1 to FF.

The default is **0x8**.

7. Configure the iSCSI priority advertised for the iSCSI protocol in Application Priority TLVs.
PROTOCOL LLDP mode

[no] iscsi priority-bits *priority-bitmap*

The priority-bitmap range is from 1 to FF.

The default is **0x10**.

DCBx Error Messages

The following syslog messages appear when an error in DCBx operation occurs.

LLDP_MULTIPLE_PEER_DETECTED: DCBx is operationally disabled after detecting more than one DCBx peer on the port interface.

LLDP_PEER_AGE_OUT: DCBx is disabled as a result of LLDP timing out on a DCBx peer interface.

DSM_DCBx_PEER_VERSION_CONFLICT: A local port expected to receive the IEEE, CIN, or CEE version in a DCBx TLV from a remote peer but received a different, conflicting DCBx version.

DSM_DCBx_PFC_PARAMETERS_MATCH and DSM_DCBx_PFC_PARAMETERS_MISMATCH: A local DCBx port received a compatible (match) or incompatible (mismatch) PFC configuration from a peer.

DSM_DCBx_ETS_PARAMETERS_MATCH and DSM_DCBx_ETS_PARAMETERS_MISMATCH: A local DCBx port received a compatible (match) or incompatible (mismatch) ETS configuration from a peer.

LLDP_UNRECOGNISED_DCBx_TLV_RECEIVED: A local DCBx port received an unrecognized DCBx TLV from a peer.

Debugging DCBx on an Interface

To enable DCBx debug traces for all or a specific control paths, use the following command.

- Enable DCBx debugging.
EXEC PRIVILEGE mode

```
debug DCBx {all | auto-detect-timer | config-exchnng | fail | mgmt | resource  
| sem | tlv}
```

- all: enables all DCBx debugging operations.
- auto-detect-timer: enables traces for DCBx auto-detect timers.
- config-exchnng: enables traces for DCBx configuration exchanges.
- fail: enables traces for DCBx failures.
- mgmt: enables traces for DCBx management frames.
- resource: enables traces for DCBx system resource frames.
- sem: enables traces for the DCBx state machine.
- tlv: enables traces for DCBx TLVs.

Verifying the DCB Configuration

To display DCB configurations, use the following `show` commands.

Table 10. Displaying DCB Configurations

Command	Output
<code>show dot1p-queue mapping</code>	Displays the current 802.1p priority-queue mapping.
<code>show dcb [stack-unit <i>unit-number</i>]</code>	Displays the data center bridging status, number of PFC-enabled ports, and number of PFC-enabled queues. On the master switch in a stack, you can specify a stack-unit number. The range is from 0 to 5.
<code>show qos priority-groups</code>	Displays the ETS priority groups configured on the switch, including the 802.1p priority classes and ID of each group.
<code>show interface <i>port-type slot/port</i> pfc {summary detail}</code>	Displays the PFC configuration applied to ingress traffic on an interface, including priorities and link delay. To clear PFC TLV counters, use the <code>clear pfc counters interface <i>port-type slot/port</i></code> command.
<code>show interface <i>port-type slot/port</i> pfc statistics</code>	Displays counters for the PFC frames received and transmitted (by dot1p priority class) on an interface.
<code>show interface <i>port-type slot/port</i> ets {summary detail}</code>	Displays the ETS configuration applied to egress traffic on an interface, including priority groups with priorities and bandwidth allocation. To clear ETS TLV counters, enter the <code>clear ets counters interface <i>port-type slot/port</i></code> command.

Example of the show dot1p-queue mapping Command

Example of the show dcb Command

Example of the show interfaces pfc summary Command

Example of the show interface pfc statistics Command

Example of the show interface ets summary Command

Example of the show interface ets detail Command

Example of the show stack-unit all stack-ports all pfc details Command

Example of the show stack-unit all stack-ports all ets details Command

Example of the show interface DCBx detail Command

```
Dell(conf)# show dot1p-queue-mapping
Dot1p Priority: 0 1 2 3 4 5 6 7
Queue          : 0 0 0 1 2 3 3 3
```

```
Dell# show dcb
stack-unit 0 port-set 0
    DCB Status : Enabled
    PFC Port Count : 56 (current), 56 (configured)
    PFC Queue Count : 2 (current), 2 (configured)
```

```
Dell# show interfaces tengigabitethernet 0/49 pfc summary
Interface TenGigabitEthernet 0/49
  Admin mode is on
  Admin is enabled
  Remote is enabled, Priority list is 4
  Remote Willing Status is enabled
  Local is enabled
  Oper status is Recommended
  PFC DCBx Oper status is Up
  State Machine Type is Feature
  TLV Tx Status is enabled
  PFC Link Delay 45556 pause quantams
  Application Priority TLV Parameters :
  -----
  FCOE TLV Tx Status is disabled
  ISCSI TLV Tx Status is disabled
  Local FCOE PriorityMap is 0x8
  Local ISCSI PriorityMap is 0x10
  Remote FCOE PriorityMap is 0x8
  Remote ISCSI PriorityMap is 0x8
```

```
Dell# show interfaces tengigabitethernet 0/49 pfc detail
Interface TenGigabitEthernet 0/49
  Admin mode is on
  Admin is enabled
  Remote is enabled
  Remote Willing Status is enabled
  Local is enabled
  Oper status is recommended
  PFC DCBx Oper status is Up
  State Machine Type is Feature
  TLV Tx Status is enabled
```



```
PFC Link Delay 45556 pause quanta
Application Priority TLV Parameters :
```

```
-----
FCOE TLV Tx Status is disabled
ISCSI TLV Tx Status is disabled
Local FCOE PriorityMap is 0x8
Local ISCSI PriorityMap is 0x10
Remote FCOE PriorityMap is 0x8
Remote ISCSI PriorityMap is 0x8
```

0 Input TLV pkts, 1 Output TLV pkts, 0 Error pkts, 0 Pause Tx pkts, 0 Pause Rx pkts

The following table describes the `show interface pfc summary` command fields.

Table 11. show interface pfc summary Command Description

Fields	Description
Interface	Interface type with stack-unit and port number.
Admin mode is on; Admin is enabled	PFC Admin mode is on or off with a list of the configured PFC priorities . When PFC admin mode is on, PFC advertisements are enabled to be sent and received from peers; received PFC configuration takes effect. The admin operational status for a DCBx exchange of PFC configuration is enabled or disabled.
Remote is enabled; Priority list Remote Willing Status is enabled	Operational status (enabled or disabled) of peer device for DCBx exchange of PFC configuration with a list of the configured PFC priorities. Willing status of peer device for DCBx exchange (Willing bit received in PFC TLV): enabled or disabled.
Local is enabled	DCBx operational status (enabled or disabled) with a list of the configured PFC priorities
Operational status (local port)	DCBx operational status (enabled or disabled) with a list of the configured PFC priorities. Port state for current operational PFC configuration: <ul style="list-style-type: none"> • Init: Local PFC configuration parameters were exchanged with peer. • Recommend: Remote PFC configuration parameters were received from peer. • Internally propagated: PFC configuration parameters were received from configuration source.
PFC DCBx Oper status	Operational status for exchange of PFC configuration on local port: match (up) or mismatch (down).
State Machine Type	Type of state machine used for DCBx exchanges of PFC parameters: <ul style="list-style-type: none"> • Feature: for legacy DCBx versions

Fields	Description
	<ul style="list-style-type: none"> Symmetric: for an IEEE version
TLV Tx Status	Status of PFC TLV advertisements: enabled or disabled.
PFC Link Delay	Link delay (in quanta) used to pause specified priority traffic.
Application Priority TLV: FCOE TLV Tx Status	Status of FCoE advertisements in application priority TLVs from local DCBx port: enabled or disabled.
Application Priority TLV: ISCSI TLV Tx Status	Status of ISCSI advertisements in application priority TLVs from local DCBx port: enabled or disabled.
Application Priority TLV: Local FCOE Priority Map	Priority bitmap used by local DCBx port in FCoE advertisements in application priority TLVs.
Application Priority TLV: Local ISCSI Priority Map	Priority bitmap used by local DCBx port in ISCSI advertisements in application priority TLVs.
Application Priority TLV: Remote FCOE Priority Map	Status of FCoE advertisements in application priority TLVs from remote peer port: enabled or disabled.
Application Priority TLV: Remote ISCSI Priority Map	Status of iSCSI advertisements in application priority TLVs from remote peer port: enabled or disabled.
PFC TLV Statistics: Input TLV pkts	Number of PFC TLVs received.
PFC TLV Statistics: Output TLV pkts	Number of PFC TLVs transmitted.
PFC TLV Statistics: Error pkts	Number of PFC error packets received.
PFC TLV Statistics: Pause Tx pkts	Number of PFC pause frames transmitted.
PFC TLV Statistics: Pause Rx pkts	Number of PFC pause frames received

```
Dell#show interfaces tengigabitethernet 0/3 pfc statistics
Interface TenGigabitEthernet 0/3
```

```
Priority Rx XOFF Frames Rx Total Frames Tx Total Frames
-----
0          0          0          0
1          0          0          0
2          0          0          0
3          0          0          0
4          0          0          0
5          0          0          0
6          0          0          0
7          0          0          0
```

```
Dell(conf)# show interfaces te 0/0 ets summary
Interface TenGigabitEthernet 0/0
Max Supported TC Groups is 4
```

Number of Traffic Classes is 8

Admin mode is on

Admin Parameters :

Admin is enabled

TC-grp	Priority#	Bandwidth	TSA
0	0,1,2,3,4,5,6,7	100%	ETS
1		0%	ETS
2		0%	ETS
3		0%	ETS
4		0%	ETS
5		0%	ETS
6		0%	ETS
7		0%	ETS

Priority#	Bandwidth	TSA
0	13%	ETS
1	13%	ETS
2	13%	ETS
3	13%	ETS
4	12%	ETS
5	12%	ETS
6	12%	ETS
7	12%	ETS

Remote Parameters:

Remote is disabled

Local Parameters :

Local is enabled

TC-grp	Priority#	Bandwidth	TSA
0	0,1,2,3,4,5,6,7	100%	ETS
1		0%	ETS
2		0%	ETS
3		0%	ETS
4		0%	ETS
5		0%	ETS
6		0%	ETS
7		0%	ETS

Priority#	Bandwidth	TSA
0	13%	ETS
1	13%	ETS
2	13%	ETS
3	13%	ETS
4	12%	ETS
5	12%	ETS
6	12%	ETS
7	12%	ETS

Oper status is init

Conf TLV Tx Status is disabled

Traffic Class TLV Tx Status is disabled

0 Input Conf TLV Pkts, 0 Output Conf TLV Pkts, 0 Error Conf TLV Pkts

0 Input Traffic Class TLV Pkts, 0 Output Traffic Class TLV Pkts, 0 Error

Traffic Class TLV

Pkts

The following table describes the show interface ets detail command fields.

Dell(conf)# show interfaces tengigabitethernet 0/0 ets detail

Interface TenGigabitEthernet 0/0

Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on
Admin Parameters :

Admin is enabled

TC-grp	Priority#	Bandwidth	TSA
0	0,1,2,3,4,5,6,7	100%	ETS
1		0%	ETS
2		0%	ETS
3		0%	ETS
4		0%	ETS
5		0%	ETS
6		0%	ETS
7		0%	ETS

Priority#	Bandwidth	TSA
0	13%	ETS
1	13%	ETS
2	13%	ETS
3	13%	ETS
4	12%	ETS
5	12%	ETS
6	12%	ETS
7	12%	ETS

Remote Parameters:

Remote is disabled

Local Parameters :

Local is enabled

TC-grp	Priority#	Bandwidth	TSA
0	0,1,2,3,4,5,6,7	100%	ETS
1		0%	ETS
2		0%	ETS
3		0%	ETS
4		0%	ETS
5		0%	ETS
6		0%	ETS
7		0%	ETS

Priority#	Bandwidth	TSA
0	13%	ETS
1	13%	ETS
2	13%	ETS
3	13%	ETS
4	12%	ETS
5	12%	ETS
6	12%	ETS
7	12%	ETS

Oper status is init

Conf TLV Tx Status is disabled

Traffic Class TLV Tx Status is disabled

0 Input Conf TLV Pkts, 0 Output Conf TLV Pkts, 0 Error Conf TLV Pkts

0 Input Traffic Class TLV Pkts, 0 Output Traffic Class TLV Pkts, 0 Error

Traffic Class TLV

Pkts

Table 12. show interface ets detail Command Description

Field	Description
Interface	Interface type with stack-unit and port number.
Max Supported TC Group	Maximum number of priority groups supported.
Number of Traffic Classes	Number of 802.1p priorities currently configured.
Admin mode	ETS mode: on or off. When on, the scheduling and bandwidth allocation configured in an ETS output policy or received in a DCBx TLV from a peer can take effect on an interface.
Admin Parameters	ETS configuration on local port, including priority groups, assigned dot1p priorities, and bandwidth allocation.
Remote Parameters	ETS configuration on remote peer port, including Admin mode (enabled if a valid TLV was received or disabled), priority groups, assigned dot1p priorities, and bandwidth allocation. If the ETS Admin mode is enabled on the remote port for DCBx exchange, the Willing bit received in ETS TLVs from the remote peer is included.
Local Parameters	ETS configuration on local port, including Admin mode (enabled when a valid TLV is received from a peer), priority groups, assigned dot1p priorities, and bandwidth allocation.
Operational status (local port)	Port state for current operational ETS configuration: <ul style="list-style-type: none"> • Init: Local ETS configuration parameters were exchanged with peer. • Recommend: Remote ETS configuration parameters were received from peer. • Internally propagated: ETS configuration parameters were received from configuration source.
ETS DCBx Oper status	Operational status of ETS configuration on local port: match or mismatch.
State Machine Type	Type of state machine used for DCBx exchanges of ETS parameters: <ul style="list-style-type: none"> • Feature: for legacy DCBx versions • Asymmetric: for an IEEE version
Conf TLV Tx Status	Status of ETS Configuration TLV advertisements: enabled or disabled.
ETS TLV Statistic: Input Conf TLV pkts	Number of ETS Configuration TLVs received.
ETS TLV Statistic: Output Conf TLV pkts	Number of ETS Configuration TLVs transmitted.

Field	Description
ETS TLV Statistic: Error Conf TLV pkts	Number of ETS Error Configuration TLVs received.

Dell(conf)# show stack-unit all stack-ports all pfc details

```
stack unit 0 stack-port all
  Admin mode is On
  Admin is enabled, Priority list is 4-5
  Local is enabled, Priority list is 4-5
  Link Delay 45556 pause quantum
  0 Pause Tx pkts, 0 Pause Rx pkts
```

```
stack unit 1 stack-port all
  Admin mode is On
  Admin is enabled, Priority list is 4-5
  Local is enabled, Priority list is 4-5
  Link Delay 45556 pause quantum
  0 Pause Tx pkts, 0 Pause Rx pkts
```

Dell(conf)# show stack-unit all stack-ports all ets details

```
Stack unit 0 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1
```

```
Admin mode is on
Admin Parameters:
```

```
-----
Admin is enabled
TC-grp      Priority#      Bandwidth      TSA
-----
```

0	0,1,2,3,4,5,6,7	100%	ETS
1		-	-
2		-	-
3		-	-
4		-	-
5		-	-
6		-	-
7		-	-
8		-	-

```
Stack unit 1 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1
```

```
Admin mode is on
Admin Parameters:
```

```
-----
Admin is enabled
TC-grp      Priority#      Bandwidth      TSA
-----
```

0	0,1,2,3,4,5,6,7	100%	ETS
1		-	-
2		-	-
3		-	-
4		-	-
5		-	-
6		-	-
7		-	-
8		-	-

```
Dell(conf)# show interface tengigabitethernet 0/49 dcbx detail
Dell#show interface te 0/49 dcbx detail

E-ETS Configuration TLV enabled           e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled          r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled           p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled   f-Application Priority for FCOE
disabled                                  disabled
I-Application priority for iSCSI enabled   i-Application Priority for iSCSI
disabled                                  disabled
```

```
-----
--
```

```
Interface TenGigabitEthernet 0/49
  Remote Mac Address 00:00:00:00:00:11
  Port Role is Auto-Upstream
  DCBX Operational Status is Enabled
  Is Configuration Source? TRUE
```

```
Local DCBX Compatibility mode is CEE
Local DCBX Configured mode is CEE
Peer Operating version is CEE
Local DCBX TLVs Transmitted: ErPfi
```

Local DCBX Status

```
-----
DCBX Operational Version is 0
DCBX Max Version Supported is 0
Sequence Number: 2
Acknowledgment Number: 2
Protocol State: In-Sync
```

Peer DCBX Status:

```
-----
DCBX Operational Version is 0
DCBX Max Version Supported is 255
Sequence Number: 2
Acknowledgment Number: 2
Total DCBX Frames transmitted 27
Total DCBX Frames received 6
Total DCBX Frame errors 0
Total DCBX Frames unrecognized 0
```

The following table describes the show interface DCBx detail command fields.

Table 13. show interface DCBx detail Command Description

Field	Description
Interface	Interface type with chassis slot and port number.
Port-Role	Configured DCBx port role: auto-upstream, auto-downstream, config-source, or manual.
DCBx Operational Status	Operational status (enabled or disabled) used to elect a configuration source and internally propagate a DCB configuration. The DCBx operational status is the combination of PFC and ETS operational status.

Field	Description
Configuration Source	Specifies whether the port serves as the DCBx configuration source on the switch: true (yes) or false (no).
Local DCBx Compatibility mode	DCBx version accepted in a DCB configuration as compatible. In auto-upstream mode, a port can only receive a DCBx version supported on the remote peer.
Local DCBx Configured mode	DCBx version configured on the port: CEE, CIN, IEEE v2.5, or Auto (port auto-configures to use the DCBx version received from a peer).
Peer Operating version	DCBx version that the peer uses to exchange DCB parameters.
Local DCBx TLVs Transmitted	Transmission status (enabled or disabled) of advertised DCB TLVs (see TLV code at the top of the show command output).
Local DCBx Status: DCBx Operational Version	DCBx version advertised in Control TLVs.
Local DCBx Status: DCBx Max Version Supported	Highest DCBx version supported in Control TLVs.
Local DCBx Status: Sequence Number	Sequence number transmitted in Control TLVs.
Local DCBx Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs.
Local DCBx Status: Protocol State	Current operational state of DCBx protocol: ACK or IN-SYNC.
Peer DCBx Status: DCBx Operational Version	DCBx version advertised in Control TLVs received from peer device.
Peer DCBx Status: DCBx Max Version Supported	Highest DCBx version supported in Control TLVs received from peer device.
Peer DCBx Status: Sequence Number	Sequence number transmitted in Control TLVs received from peer device.
Peer DCBx Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs received from peer device.
Total DCBx Frames transmitted	Number of DCBx frames sent from local port.
Total DCBx Frames received	Number of DCBx frames received from remote peer port.
Total DCBx Frame errors	Number of DCBx frames with errors received.
Total DCBx Frames unrecognized	Number of unrecognizable DCBx frames received.


QoS dot1p Traffic Classification and Queue Assignment

The following section describes QoS dot1P traffic classification and assignments.

DCB supports PFC, ETS, and DCBx to handle converged Ethernet traffic that is assigned to an egress queue according to the following QoS methods:


Honor dot1p You can honor dot1p priorities in ingress traffic at the port or global switch level (refer to Default dot1p to Queue Mapping) using the `service-class dynamic dot1p` command in INTERFACE configuration mode (refer to [Honoring dot1p Values on Ingress Packets](#)).

Layer 2 class maps You can use dot1p priorities to classify traffic in a class map and apply a service policy to an ingress port to map traffic to egress queues (refer to [Policy-Based QoS Configurations](#)).

 **NOTE:** Dell Networking does not recommend mapping all ingress traffic to a single queue when using PFC and ETS. However, Dell Networking does recommend using Ingress traffic classification using the `service-class dynamic dot1p` command (honor dot1p) on all DCB-enabled interfaces. If you use L2 class maps to map dot1p priority traffic to egress queues, take into account the default dot1p-queue assignments in the following table and the maximum number of two lossless queues supported on a port (refer to [Configuring Lossless Queues](#)).

Although the system allows you to change the default dot1p priority-queue assignments (refer to [Setting dot1p Priorities for Incoming Traffic](#)), DCB policies applied to an interface may become invalid if you reconfigure dot1p-queue mapping. If the configured DCB policy remains valid, the change in the dot1p-queue assignment is allowed. For DCB ETS enabled interfaces, traffic destined to queue that is not mapped to any dot1p priority are dropped.

dot1p Value in the Incoming Frame	Egress Queue Assignment
0	0
1	0
2	0
3	1
4	2
5	3
6	3
7	3

 **NOTE:** If you reconfigure the global dot1p-queue mapping, an automatic re-election of the DCBX configuration source port is performed (refer to [Configuration Source Election](#)).

Configuring the Dynamic Buffer Method

To configure the dynamic buffer capability, perform the following steps:

1. Enable the DCB application. By default, DCB is enabled and link-level flow control is disabled on all interfaces.

```
CONFIGURATION mode
```

```
S6000-109-Dell(conf)#dcb enable
```

2. Configure the shared PFC buffer size and the total buffer size. A maximum of 4 lossless queues are supported.

```
CONFIGURATION mode
```

```
S6000-109-Dell(conf)#dcb pfc-shared-buffer-size 4000
```

```
S6000-109-Dell(conf)#dcb pfc-total-buffer-size 5000
```



NOTE: For `dcb pfc-shared-buffer-size`, the range is from `<0-11210>` in KB (default LC=2496/SFM=3328)

For `dcb pfc-total-buffer-size`, the range is from `<0-11210>` in KB (default LC=7488/SFM=7596)

3. Configure the number of PFC queues.

```
CONFIGURATION mode
```

```
Dell(conf)#dcb enable pfc-queues 4
```

The number of ports supported based on lossless queues configured will depend on the buffer.

For each priority, you can specify the shared buffer threshold limit, the ingress buffer size, buffer limit for pausing the acceptance of packets, and the buffer offset limit for resuming the acceptance of received packets.

4. Configure the profile name for the DCB buffer threshold

```
CONFIGURATION mode
```

```
Dell(conf)#dcb-buffer-threshold test
```

5. DCB-BUFFER-THRESHOLD mode

```
Dell(conf-dcb-buffer-thr)# priority 0 buffer-size 52 pause-threshold 16  
resume-offset 10 shared-threshold-weight 7
```

6. Assign the DCB policy to the DCB buffer threshold profile on stack ports.

```
CONFIGURATION mode
```

```
Dell(conf)# dcb-policy buffer-threshold stack-unit all stack-ports all test
```

7. Assign the DCB policy to the DCB buffer threshold profile on interfaces. This setting takes precedence over the default buffer-threshold setting.

```
INTERFACE mode (conf-if-te)
```

```
Dell(conf-if-te-0/0)#dcb-policy buffer-threshold test
```

8. Create a QoS policy buffer and enter the QoS Policy Buffer Configuration mode to configure the no-drop queues, ingress buffer size, buffer limit for pausing, and buffer offset limit for resuming.
CONFIGURATION mode

```
Dell(conf)# qos-policy-buffer test
```

```
Dell (conf-qos-policy-buffer)#queue 0 pause no-drop buffer-size 128000  
pause-threshold 103360 resume-threshold 83520
```

```
Dell (conf-qos-policy-buffer)# queue 4 pause no-drop buffer-size 128000  
pause-threshold 103360 resume-threshold 83520
```

Debugging and Diagnostics


This chapter describes debugging and diagnostics for the MXL switch.

Offline Diagnostics

The offline diagnostics test suite is useful for isolating faults and debugging hardware.

The diagnostics tests are grouped into three levels:

- **Level 0** — Level 0 diagnostics check for the presence of various components and perform essential path verifications. In addition, Level 0 diagnostics verify the identification registers of the components on the board.
- **Level 1** — A smaller set of diagnostic tests. Level 1 diagnostics perform status, self-test, access, and read-write tests for all the components on the board and test their registers for appropriate values. In addition, Level 1 diagnostics perform extensive tests on memory devices (for example, SDRAM, flash, NVRAM, EEPROM) wherever possible.
- **Level 2** — The full set of diagnostic tests. Level 2 diagnostics are used primarily for on-board MAC level, Physical level, external Loopback tests, and more extensive component diagnostics. Various components on the board are put into Loopback mode and test packets are transmitted through those components. These diagnostics also perform snake tests using virtual local area network (VLAN) configurations.

 **NOTE:** Diagnostic is not allowed in Stacking mode, including member stacking. Avoid stacking before executing the diagnostic tests in the chassis.

Important Points to Remember

- You can only perform offline diagnostics on an offline standalone unit. You cannot perform diagnostics if the ports are configured in a stacking group. Remove the port(s) from the stacking group before executing the diagnostic test.
- Diagnostics only test connectivity, not the entire data path.
- Diagnostic results are stored on the flash of the unit on which you performed the diagnostics.
- When offline diagnostics are complete, the unit or stack member reboots automatically.

Running Offline Diagnostics


To run offline diagnostics, use the following commands.
For more information, refer to the examples following the steps.

1. Place the unit in the offline state.

EXEC Privilege mode

```
offline stack-unit <id>
```

You cannot enter this command on a MASTER or Standby stack unit.

 **NOTE:** The system reboots when the offline diagnostics complete. This is an automatic process. The following warning message appears when you implement the `offline stack-unit <id>` command: Warning - Diagnostic execution will cause stack-unit to reboot after completion of diags. Proceed with Offline-Diags [confirm yes/no]:y

```
Dell#offline stack-unit 0
Warning - offline of unit will bring down all the protocols and
the unit will be operationally down, except for running Diagnostics.
Please make sure that stacking/fanout not configured for Diagnostics
execution.
Also reboot/online command is necessary for normal operation after the
offline command is issued.
Proceed with Offline [confirm yes/no]:yes
Dell#Dec 15 03:58:37: %STKUNIT0-M:CP %CHMGR-2-STACKUNIT_DOWN: Stack unit 0
down - stack unit offline
```

2. Confirm the offline status.

EXEC Privilege mode

show system brief

```
Dell#show system brief
```

```
Stack MAC : 00:1e:c9:f1:00:cb
Reload-Type : normal-reload [Next boot : normal-reload]
```

```
-- Stack Info --
```

Unit	UnitType	Status	ReqTyp	CurTyp	Version	Ports
0	Management	offline	MXL-10/40GbE	MXL-10/40GbE	9.4	(0.0)
56						
1	Member	not present				
2	Member	not present				
3	Member	not present				
4	Member	not present				
5	Member	not present				

```
Dell#
```

3. Start diagnostics on the unit.diag

When the tests are complete, the system displays the following message and reboots the unit automatically

```
Diags completed... Rebooting the system now!!!
Dec 15 04:00:38: %MXL-10/40GbE:0 %DIAGAGT-6-DA_DIAG_DONE: Diags finished on
stack unit 0
```

Diagnostic results are printed to a file in the flash using the filename format `TestReport-SU-stack-unit.txt`. Log messages differ somewhat when diagnostics are done on a standalone unit and on a stack member.

Example of the diag command (Standalone unit)

```
Dell#diag stack-unit 0 level0
Warning - diagnostic execution will cause multiple link flaps on the peer side
- advisable to shut directly connected ports
Proceed with Diags [confirm yes/no]: yes
FTOS#Dec 15 04:14:07: %MXL-10/40GbE:0 %DIAGAGT-6-DA_DIAG_STARTED: Starting
diags on stack unit 0
00:12:10 : System may take additional time for Driver Init.
00:12:10 : Approximate time to complete the Diags ... 6 Mins

00:13:53 : Diagnostic test results are stored on file: flash:/TestReport-
SU-0.txt
Diags completed... Rebooting the system now!!!
Dec 15 04:15:54: %MXL-10/40GbE:0 %DIAGAGT-6-DA_DIAG_DONE: Diags finished on
stack unit 0
syncing disks... 1 1 done
unmounting file systems...
unmounting /f10/flash (/dev/ld0e)...
unmounting /usr/pkg (/dev/ld0h)...
unmounting /usr (mfs:35)...
unmounting /lib (mfs:24)...
unmounting /f10 (mfs:21)...
unmounting /tmp (mfs:15)...
unmounting /kern (kernfs)...
unmounting / (/dev/md0a)... done
rebooting...
```

Example of the show file flash:\\ command (Standalone Unit)

```
Dell#show file flash://TestReport-SU-0.txt

*****BLADE IOM
DIAGNOSTICS*****

          Board                : Blade IOM Dell Inc.
          CPU Version           : XLP3XX-A0
          Stack Unit Board Temp : 49 Degree C
          Stack Unit Number     : 0
          Board Serial Number   : TW282981C80067
          Board Type            : Blade IOM Module
          CPLD Revision         : 0x6
          Image Build Version   : 9-4(0-89)

***** BLADE IOM LEVEL 0
DIAGNOSTICS*****

Test 1 - Power Rail Status Test ..... PASS
Test 2.000 - OptMod: Power Status Test ..... PASS
Test 2.001 - OptMod: Power Status Test ..... PASS
Test 2 - OptMod: Power Status Test ..... PASS
Test 3.000 - Board Temperature Sensor Test ..... PASS
Test 3.001 - Board Temperature Sensor Test ..... PASS
Test 3 - Board Temperature Sensor Test ..... PASS
Test 4 - RTC Presence Test ..... PASS
Test 5.000 - CPU Sdram Presence Test ..... PASS
Test 6.000 - CPU Sdram Size Test ..... PASS
diagBladeIOMUsbAAccessTest[238]: ERROR: No USB A device found
Test 7 - USB A Access Test ..... NOT
PRESENT
diagBladeIOMUsbAPresenceGet[267]: ERROR: No USB device found
diagBladeIOMUsbHostControllerAccessTest[608]: ERROR: No USB device detected.
Test 8 - Usb Host Controller Access Test ..... NOT
```

```

PRESENT
Test 9 - SD Flash Access Test ..... PASS
Test 10.000 - Qsfp Plus Power Mode Test ..... PASS
Test 10.001 - Qsfp Plus Power Mode Test ..... PASS
Test 10 - Qsfp Plus Power Mode Test ..... PASS
Test 11 - CPLD Presence Test ..... PASS
Test 12 - Flash Access Test ..... PASS
Test 13 - Board Revision Test ..... PASS
Test 14 - MGMT PHY Presence Test ..... PASS
Test 15.000 - Optional Module Type Test ..... PASS
Test 15.001 - Optional Module Type Test ..... PASS
Test 15 - Optional Module Type Test ..... PASS
Test 16.000 - Qsfp Plus Presence Test ..... PASS
Test 16.001 - Qsfp Plus Presence Test ..... PASS
Test 16 - Qsfp Plus Presence Test ..... PASS
Test 17 - Cpu Type Detect Test ..... PASS

```

```

***** BLADE IOM LEVEL 1
DIAGNOSTICS*****

```

```

Test 101 - RTC Function Test ..... PASS
Test 102 - RTC Rollover Test ..... PASS
Test 103 - GPIO Access Test ..... PASS
Test 104 - PSoC Access Test ..... PASS
Test 105 - PCIe BCM56846 Access Test ..... PASS
Test 106 - CPU SDRAM Access Test ..... PASS
Test 107 - CPU SDRAM Data Line Test ..... PASS
Test 108 - CPU SDRAM Address Line Test ..... PASS
diagBladeIOMUsbAPresenceGet[267]: ERROR: No USB device found
diagBladeIOMUsbFileCopyTest[92]: ERROR: No USB device detected.
Test 109 - Usb File Copy Stress Test ..... NOT
PRESENT
Test 110 - Flash Rw Test ..... PASS
Test 111 - I2C Stress Test ..... PASS

```

Trace Logs

In addition to the syslog buffer, the Dell Networking OS buffers trace messages which are continuously written by various software tasks to report hardware and software events and status information.

Each trace message provides the date, time, and name of the Dell Networking OS process. All messages are stored in a ring buffer. You can save the messages to a file either manually or automatically after failover.

Auto Save on Crash or Rollover

Exception information for MASTER or standby units is stored in the *flash://TRACE_LOG_DIR* directory. This directory contains files that save trace information when there has been a task crash or timeout.

- On a MASTER unit, you can reach the *TRACE_LOG_DIR* files by FTP or by using the `show file` command from the *flash://TRACE_LOG_DIR* directory.
- On a Standby unit, you can reach the *TRACE_LOG_DIR* files only by using the `show file` command from the *flash://TRACE_LOG_DIR* directory.

 **NOTE:** Non-management member units do not support this functionality.


Example of the `dir flash:` Command

```
Dell#dir flash://TRACE_LOG_DIR
Directory of flash://TRACE_LOG_DIR
 1 drwx   4096 Jan 17 2011 15:02:16 +00:00 .
 2 drwx   4096 Jan 01 1980 00:00:00 +00:00 ..
 3 -rwx 100583 Feb 11 2011 20:41:36 +00:00 failure_trace0_RPM0_CP

flash: 2143281152 bytes total (2069291008 bytes free)
```

Using the Show Hardware Commands

The `show hardware` command tree consists of commands used with the MXL switch. These commands display information from a hardware sub-component and from hardware-based feature tables.

 **NOTE:** Use the `show hardware` commands only under the guidance of the Dell Technical Assistance Center.

- View internal interface status of the stack-unit CPU port which connects to the external management interface.

EXEC Privilege mode

```
show hardware stack-unit {0-5} cpu management statistics
```

- View driver-level statistics for the data-plane port on the CPU for the specified stack-unit.

EXEC Privilege mode

```
show hardware stack-unit {0-5} cpu data-plane statistics
```

This view provides insight into the packet types entering the CPU to see whether CPU-bound traffic is internal (IPC traffic) or network control traffic, which the CPU must process.

- View the modular packet buffers details per stack unit and the mode of allocation.

EXEC Privilege mode

```
show hardware stack-unit {0-5} buffer total-buffer
```

- View the modular packet buffers details per unit and the mode of allocation.

EXEC Privilege mode

```
show hardware stack-unit {0-5} buffer unit {0-1} total-buffer
```

- View the forwarding plane statistics containing the packet buffer usage per port per stack unit.

EXEC Privilege mode

```
show hardware stack-unit {0-5} buffer unit {0-1} port {1-64 | all} buffer-
info
```

- View the forwarding plane statistics containing the packet buffer statistics per COS per port.

EXEC Privilege mode

```
show hardware stack-unit {0-5} buffer unit {0-1} port {1-64} queue {0-14 |
all} buffer-info
```

- View input and output statistics on the party bus, which carries inter-process communication traffic between CPUs.

EXEC Privilege mode


```
show hardware stack-unit {0-5} cpu party-bus statistics
```

- View the ingress and egress internal packet-drop counters, MAC counters drop, and FP packet drops for the stack unit on per port basis.

EXEC Privilege mode

```
show hardware stack-unit {0-5} drops unit {0-0} port {33-56}
```

This view helps identifying the stack unit/port pipe/port that may experience internal drops.

- View the input and output statistics for a stack-port interface.

EXEC Privilege mode

```
show hardware stack-unit {0-5} stack-port {33-56}
```

- View the counters in the field processors of the stack unit.

EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} counters
```

- View the details of the FP Devices and Hi gig ports on the stack-unit.

EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} details
```

- Execute a specified bShell command from the CLI without going into the bShell.

EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} execute-shell-cmd {command}
```

- View the Multicast IPMC replication table from the bShell.

EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} ipmc-replication
```

- View the internal statistics for each port-pipe (unit) on per port basis.

EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} port-stats [detail]
```

- View the stack-unit internal registers for each port-pipe.

EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} register
```

- View the tables from the bShell through the CLI without going into the bShell.

EXEC Privilege mode

```
show hardware stack-unit {0-5} unit {0-0} table-dump {table name}
```

Enabling Environmental Monitoring

The MXL switch components use environmental monitoring hardware to detect transmit power readings, receive power readings, and temperature updates.

To receive periodic power updates, you must enable the following command.

- Enable environmental monitoring.

```
enable optic-info-update interval
```

Example of the show interfaces transceiver Command

```
Dell#show int ten 0/49 transceiver
SFP is present
SFP 49 Serial Base ID fields
SFP 49 Id = 0x03
SFP 49 Ext Id = 0x04
SFP 49 Connector = 0x07
SFP 49 Transceiver Code = 0x00 0x00 0x00 0x01 0x20 0x40 0x0c 0x01
SFP 49 Encoding = 0x01
SFP 49 BR Nominal = 0x0c
SFP 49 Length(9um) Km = 0x00
SFP 49 Length(9um) 100m = 0x00
SFP 49 Length(50um) 10m = 0x37
SFP 49 Length(62.5um) 10m = 0x1e
SFP 49 Length(Copper) 10m = 0x00
SFP 49 Vendor Rev =
SFP 49 Laser Wavelength = 850 nm
SFP 49 CheckCodeBase = 0x78
SFP 49 Serial Extended ID fields
SFP 49 Options = 0x00 0x12
SFP 49 BR max = 0
SFP 49 BR min = 0
SFP 49 Vendor SN = P11C0B0
SFP 49 Datecode = 020919
SFP 49 CheckCodeExt = 0xb6

SFP 49 Diagnostic Information
=====
SFP 49 Rx Power measurement type = Average
=====
SFP 49 Temp High Alarm threshold = 100.000C
SFP 49 Voltage High Alarm threshold = 5.000V
SFP 49 Bias High Alarm threshold = 100.000mA
SFP 49 TX Power High Alarm threshold = 5.000mW
SFP 49 RX Power High Alarm threshold = 5.000mW
SFP 49 Temp Low Alarm threshold = -50.000C
SFP 49 Voltage Low Alarm threshold = 0.000V
SFP 49 Bias Low Alarm threshold = 0.000mA
SFP 49 TX Power Low Alarm threshold = 0.000mW
SFP 49 RX Power Low Alarm threshold = 0.000mW
=====
SFP 49 Temp High Warning threshold = 100.000C
SFP 49 Voltage High Warning threshold = 5.000V
SFP 49 Bias High Warning threshold = 100.000mA
SFP 49 TX Power High Warning threshold = 5.000mW
SFP 49 RX Power High Warning threshold = 5.000mW
SFP 49 Temp Low Warning threshold = -50.000C
SFP 49 Voltage Low Warning threshold = 0.000V
SFP 49 Bias Low Warning threshold = 0.000mA
SFP 49 TX Power Low Warning threshold = 0.000mW
SFP 49 RX Power Low Warning threshold = 0.000mW
=====
SFP 49 Temperature = 40.844C
SFP 49 Voltage = 3.169V
SFP 49 Tx Bias Current = 0.000mA
SFP 49 Tx Power = 0.000mW
SFP 49 Rx Power = 0.227mW
=====
SFP 49 Data Ready state Bar = False
SFP 49 Rx LOS state = False
SFP 49 Tx Fault state = False
```

Recognize an Over-Temperature Condition

An overtemperature condition occurs, for one of two reasons: the card genuinely is too hot or a sensor has malfunctioned.

Inspect cards adjacent to the one reporting the condition to discover the cause.

- If directly adjacent cards are not normal temperature, suspect a genuine overheating condition.
- If directly adjacent cards are normal temperature, suspect a faulty sensor.

When the system detects a genuine over-temperature condition, it powers off the card. To recognize this condition, look for the following system messages:

```
CHMGR-2-MAJOR_TEMP: Major alarm: chassis temperature high (temperature reaches
or exceeds threshold of
[value]C)
CHMGR-2-TEMP_SHUTDOWN_WARN: WARNING! temperature is [value]C; approaching
shutdown threshold of [value]C
```

To view the programmed alarm thresholds levels, including the shutdown value, use the `show alarms threshold` command.

Example of the `show alarms threshold` Command

```
Dell#show alarms threshold


-- Temperature Limits (deg C) --
-----
      BelowNormal Normal Elevated Critical Trip/Shutdown
Unit0 <=40      41      71      81      86
Dell#
```

Troubleshoot an Over-Temperature Condition

To troubleshoot an over-temperature condition, use the following information.

1. Use the `show environment` commands to monitor the temperature levels.
2. Check air flow through the system. Ensure that the air ducts are clean and that all fans are working correctly.
3. After the software has determined that the temperature levels are within normal limits, you can re-power the card safely. To bring back the line card online, use the `power-on` command in EXEC mode.

In addition, Dell Networking requires that you install blanks in all slots without a line card to control airflow for adequate system cooling.

 **NOTE:** Exercise care when removing a card; if it has exceeded the major or shutdown thresholds, the card could be hot to the touch.

Example of the `show environment` Command

```
Dell#show environment

-- Unit Environment Status --
Unit Status Temp Voltage
-----
* 0 online 71C ok
```

* Management Unit

-- Thermal Sensor Readings (deg C) --

Unit Sensor0 Sensor1 Sensor2 Sensor3 Sensor4 Sensor5 Sensor6 Sensor7 Sensor8
Sensor9

0 45 43 66 61 66 62 70 65 67 71

Recognize an Under-Voltage Condition

If the system detects an under-voltage condition, it sends an alarm.

To recognize this condition, look for the following system message: %CHMGR-1-CARD_SHUTDOWN:

Major alarm: Line card 2 down - auto-shutdown due to under voltage.


This message indicates that the specified card is not receiving enough power. In response, the system first shuts down Power over Ethernet (PoE).

Troubleshoot an Under-Voltage Condition

To troubleshoot an under-voltage condition, check that the correct number of power supplies are installed and their Status light emitting diodes (LEDs) are lit.

The following table lists information for SNMP traps and OIDs on S-Series environmental monitoring hardware and hardware components.

Table 14. SNMP Traps and OIDs

OID String	OID Name	Description
Receiving Power		
.1.3.6.1.4.1.6027.3.10.1.2.5.1.6	chSysPortXfpRecvPower	OID displays the receiving power of the connected optics.
Transmitting power		
.1.3.6.1.4.1.6027.3.10.1.2.5.1.8	chSysPortXfpTxPower	OID displays the transmitting power of the connected optics.
Temperature		
.1.3.6.1.4.1.6027.3.10.1.2.5.1.7	chSysPortXfpRecvTemp	OID displays the temperature of the connected optics.
 NOTE: These OIDs only generate if you enable the <code>enable optic-info-update-interval</code> is enabled command.		
Hardware MIB Buffer Statistics		
.1.3.6.1.4.1.6027.3.16.1.1.4	fpPacketBufferTable	View the modular packet buffers details per stack unit and the mode of allocation.

OID String	OID Name	Description
.1.3.6.1.4.1.6027.3.16.1.1.5	fpStatsPerPortTable	View the forwarding plane statistics containing the packet buffer usage per port per stack unit.
.1.3.6.1.4.1.6027.3.16.1.1.6	fpStatsPerCOSTable	View the forwarding plane statistics containing the packet buffer statistics per COS per port.

Troubleshooting Packet Loss

The `show hardware stack-unit` command is intended primarily to troubleshoot packet loss. To troubleshoot packet loss, use the following commands.

- `show hardware stack-unit 0-5 cpu data-plane statistics`
- `show hardware stack-unit 0-5 cpu party-bus statistics`
- `show hardware stack-unit 0-5 drops unit 0-0 port 1-56`
- `show hardware stack-unit 0-5 stack-port 33-56`
- `show hardware stack-unit 0-5 unit 0-0 {counters | details | port-stats [detail] | register | ipmc-replication | table-dump}:`
- `show hardware {layer2| layer3} {eg acl |in acl} stack-unit 0-5 port-set 0-0`
- `show hardware layer3 qos stack-unit 0-5 port-set 0-0`
- `show hardware system-flow layer2 stack-unit 0-5 port-set 0-1 [counters]`
- `clear hardware stack-unit 0-5 counters`
- `clear hardware stack-unit 0-5 cpu data-plane statistics`
- `clear hardware stack-unit 0-5 cpu party-bus statistics`
- `clear hardware stack-unit 0-5 stack-port 33-56`

Displaying Drop Counters

To display drop counters, use the following commands.

- Identify which stack unit, port pipe, and port is experiencing internal drops.
`show hardware stack-unit 0-11 drops [unit 0 [port 0-63]]`
- Display drop counters.
`show hardware stack-unit drops unit port`
- Identify which interface is experiencing internal drops.
`show hardware drops interface interface`

Example of the `show hardware stack-unit` Command to View Drop Counters Statistics

```
Dell#show hardware stack-unit 0 drops
UNIT No: 0
Total Ingress Drops :0
Total IngMac Drops :0
Total Mmu Drops :0
Total EgMac Drops :0
Total Egress Drops :0
UNIT No: 1
Total Ingress Drops :0
```

```
Total IngMac Drops :0
Total Mmu Drops :0
Total EgMac Drops :0
Total Egress Drops :0
```

```
Dell#show hardware stack-unit 0 drops unit 0
Port# :Ingress Drops :IngMac Drops :Total Mmu Drops :EgMac Drops :Egress
Drops
1 0 0 0 0 0
2 0 0 0 0 0
3 0 0 0 0 0
4 0 0 0 0 0
5 0 0 0 0 0
6 0 0 0 0 0
7 0 0 0 0 0
8 0 0 0 0 0
```

Dataplane Statistics

The `show hardware stack-unit cpu data-plane statistics` command provides insight into the packet types coming to the CPU.

The command output in the following example has been augmented, providing detailed RX/ TX packet statistics on a per-queue basis. The objective is to see whether CPU-bound traffic is internal (so-called *party bus* or IPC traffic) or network control traffic, which the CPU must process.

Example of Viewing Dataplane Statistics

Example of Viewing Party Bus Statistics

```
Dell#show hardware stack-unit 2 cpu data-plane statistics
```

```
bc pci driver statistics for device:
rxHandle          :0
noMhdr            :0
noMbuf            :0
noClus            :0
recvd             :0
dropped           :0
recvToNet         :0
rxError           :0
rxDatapathErr    :0
rxPkt (COS0)     :0
rxPkt (COS1)     :0
rxPkt (COS2)     :0
rxPkt (COS3)     :0
rxPkt (COS4)     :0
rxPkt (COS5)     :0
rxPkt (COS6)     :0
rxPkt (COS7)     :0
rxPkt (UNIT0)    :0
rxPkt (UNIT1)    :0
rxPkt (UNIT2)    :0
rxPkt (UNIT3)    :0
transmitted       :0
txRequested       :0
noTxDesc          :0
txError           :0
txReqTooLarge     :0
txInternalError   :0
txDatapathErr    :0
txPkt (COS0)     :0
```

```

txPkt (COS1)      :0
txPkt (COS2)      :0
txPkt (COS3)      :0
txPkt (COS4)      :0
txPkt (COS5)      :0
txPkt (COS6)      :0
txPkt (COS7)      :0
txPkt (UNIT0)     :0

```

The `show hardware stack-unit cpu party-bus statistics` command displays input and output statistics on the party bus, which carries inter-process communication traffic between CPUs

```

Dell#sh hardware stack-unit 2 cpu party-bus statistics
Input Statistics:
  27550 packets, 2559298 bytes
  0 dropped, 0 errors
Output Statistics:
  1649566 packets, 1935316203 bytes
  0 errors

```

Display Stack Port Statistics

The `show hardware stack-unit stack-port` command displays input and output statistics for a stack-port interface.

Example of Viewing Stack Unit Statistics

```

Dell#show hardware stack-unit 2 stack-port 49
Input Statistics:
  27629 packets, 3411731 bytes
  0 64-byte pkts, 27271 over 64-byte pkts, 207 over 127-byte pkts
  17 over 255-byte pkts, 56 over 511-byte pkts, 78 over 1023-byte pkts
  0 Multicasts, 5 Broadcasts
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  1649714 packets, 1948622676 bytes, 0 underruns
  0 64-byte pkts, 27234 over 64-byte pkts, 107970 over 127-byte pkts
  34 over 255-byte pkts, 504838 over 511-byte pkts, 1009638 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts, 1649714 Unicasts
  0 throttles, 0 discarded, 0 collisions
Rate info (interval 45 seconds):
  Input 00.00 Mbits/sec,    2 packets/sec, 0.00% of line-rate
  Output 00.06 Mbits/sec,   8 packets/sec, 0.00% of line-rate
Dell#

```

Displaying Stack Member Counters

The `show hardware stack-unit 0-5 {counters | details | port-stats [detail] | register}` command displays internal receive and transmit statistics, based on the selected command option.

The following example is a sample of the output for the `counters` option.

Example of Displaying Stack Unit Counters

Example of Displaying Counter Information for a Specific Interface

```

RIPC4.ge0      :   1,202      +1,202
RUC.ge0        :   1,224      +1,217

```

RDBG0.ge0	:	34	+24
RDBG1.ge0	:	366	+235
RDBG5.ge0	:	16	+12
RDBG7.ge0	:	18	+12
GR64.ge0	:	5,176	+24
GR127.ge0	:	1,566	+1,433
GR255.ge0	:	4	+4
GRPKT.ge0	:	1,602	+1,461
GRBYT.ge0	:	117,600	+106,202
GRMCA.ge0	:	366	+235
GRBCA.ge0	:	12	+9
GT64.ge0	:	4	+3
GT127.ge0	:	964	+964
GT255.ge0	:	4	+4
GT511.ge0	:	1	+1
GTPKT.ge0	:	973	+972
GTBCA.ge0	:	1	+1
GTBYT.ge0	:	71,531	+71,467
RUC.cpu0	:	972	+971
TDBG6.cpu0	:	1,584	+1,449=

Enabling Application Core Dumps

Application core dumps are disabled by default.

A core dump file can be very large. Due to memory requirements the file can only be sent directly to an FTP server; it is not stored on the local flash.

To enable full application core dumps, use the following command.

- Enable RPM core dumps and specify the Shutdown mode.

CONFIGURATION mode

```
logging coredump server
```

To undo this command, use the `no logging coredump server` command.

Mini Core Dumps

The Dell Networking OS supports mini core dumps on the application and kernel crashes. The mini core dump applies to Master, Standby, and Member units.

Application and kernel mini core dumps are always enabled. The mini core dumps contain the stack space and some other minimal information that you can use to debug a crash. These files are small files and are written into flash until space is exhausted. When the flash is full, the write process is stopped.

A mini core dump contains critical information in the event of a crash. Mini core dump files are located in *flash:/ (root dir)*. The application mini core filename format is *f10StkUnit<Stack_unit_no>.<Application name>.acore.mini.txt*. The kernel mini core filename format is *f10StkUnit<Stack_unit_no>.kcore.mini.txt*. The following are sample filenames.

When a member or standby unit crashes, the mini core file gets uploaded to master unit. When the master unit crashes, the mini core file is uploaded to new master.

In the MXL Switch, only the master unit has the ability to upload the coredump.

The panic string contains key information regarding the crash. Several panic string types exist, and they are displayed in regular English text to allow easier understanding of the crash cause.

Example of Application Mini Core Dump Listings

Example of a Mini Core Text File

```
Dell#dir
Directory of flash:

 1 drw-  16384 Jan  01 1980 00:00:00 +00:00 .
 2 drwx  1536 Sep  03 2009 16:51:02 +00:00 ..
 3 drw-  512 Aug  07 2009 13:05:58 +00:00 TRACE_LOG_DIR
 4 d---  512 Aug  07 2009 13:06:00 +00:00 ADMIN_DIR
 5 -rw-  8693 Sep  03 2009 16:50:56 +00:00 startup-config
 6 -rw-  8693 Sep  03 2009 16:44:22 +00:00 startup-config.bak
 7 -rw-  156 Aug  28 2009 16:16:10 +00:00 f10StkUnit0.mrtm.acore.mini.txt
 8 -rw-  156 Aug  28 2009 17:17:24 +00:00 f10StkUnit0.vrrp.acore.mini.txt
 9 -rw-  156 Aug  28 2009 18:25:18 +00:00 f10StkUnit0.sysd.acore.mini.txt
10 -rw-  156 Aug  28 2009 19:07:36 +00:00 f10StkUnit0.frrp.acore.mini.txt
11 -rw-  156 Aug  31 2009 16:18:50 +00:00 f10StkUnit2.sysd.acore.mini.txt
12 -rw-  156 Aug  29 2009 14:28:34 +00:00 f10StkUnit0.ipml.acore.mini.txt
13 -rw-  156 Aug  31 2009 16:14:56 +00:00 f10StkUnit0.acl.acore.mini.txt

flash: 3104256 bytes total (2959872 bytes free)
Dell#
```

```
VALID MAGIC
-----PANIC STRING -----
panic string is :<null>
-----STACK TRACE START-----
0035d60c <f10_save_mmu+0x120>:
00274f8c <panic+0x144>:
0024e2b0 <db_fncall+0x134>:
0024dee8 <db_command+0x258>:
0024d9c4 <db_command_loop+0xc4>:
002522b0 <db_trap+0x158>:
0026a8d0 <mi_switch+0x1b0>:
0026a00c <bpendtsleep>:
-----STACK TRACE END-----

-----FREE MEMORY-----
uvmexp.free = 0x2312
```

Enabling TCP Dumps

A TCP dump captures CPU-bound control plane traffic to improve troubleshooting and system manageability. When you enable TCP dump, it captures all the packets on the local CPU, as specified in the CLI.

You can save the traffic capture files to flash, FTP, SCP, or TFTP. The files saved on the flash are located in the `flash://TCP_DUMP_DIR/Tcpdump_<time_stamp_dir>/` directory and labeled `tcpdump_*.pcap`. There can be up to 20 `Tcpdump_<time_stamp_dir>` directories. The 21st file overwrites the oldest saved file. The maximum file size for a TCP dump capture is 1MB. When a file reaches 1MB, a new file is created, up to the specified total number of files.

Maximize the number of packets recorded in a file by specifying the `snap-length` to capture the file headers only.

The `tcpdump` command has a finite run process. When you enable the `tcpdump` command, it runs until the capture-duration timer and/or the packet-count counter threshold is met. If you do not set a threshold, the system uses a default of a 5 minute capture-duration and/or a single 1k file as the stopping point for the dump.

You can use the capture-duration timer and the packet-count counter at the same time. The TCP dump stops when the first of the thresholds is met. That means that even if the duration timer is 9000 seconds, if the maximum file count parameter is met first, the dumps stop.

To enable a TCP dump, use the following command.

- Enable a TCP dump for CPU bound traffic.

CONFIGURATION mode

```
tcpdump cp [capture-duration time | filter expression | max-file-count value
| packet-count value | snap-length value | write-to path]
```

Enabling Buffer Statistics Tracking

You can enable the tracking of statistical values of buffer spaces at a global level. The buffer statistics tracking utility operates in the max use count mode that enables the collection of maximum values of counters.

To configure the buffer statistics tracking utility, perform the following step:

Enable the buffer statistics tracking utility and enter the Buffer Statistics Snapshot configuration mode.

CONFIGURATION mode

```
Dell(conf)#buffer-stats-snapshot
```

```
Dell(conf)#no disable
```

You must enable this utility to be able to configure the parameters for buffer statistics tracking. By default, buffer statistics tracking is disabled.

Dynamic Host Configuration Protocol (DHCP)

The dynamic host configuration protocol (DHCP) is an application layer protocol that dynamically assigns IP addresses and other configuration parameters to network end-stations (hosts) based on configuration policies determined by network administrators.

DHCP relieves network administrators of manually configuring hosts, which can be a tedious and error-prone process when hosts often join, leave, and change locations on the network and it reclaims IP addresses that are no longer in use to prevent address exhaustion.

DHCP is based on a client-server model. A host discovers the DHCP server and requests an IP address, and the server either leases or permanently assigns one. There are three types of devices that are involved in DHCP negotiation:

DHCP Server	This is a network device offering configuration parameters to the client.
DHCP Client	This is a network device requesting configuration parameters from the server.
Relay Agent	This is an intermediary network device that passes DHCP messages between the client and server when the server is not on the same subnet as the host.

DHCP Packet Format and Options

DHCP uses the user datagram protocol (UDP) as its transport protocol.

The server listens on port 67 and transmits to port 68; the client listens on port 68 and transmits to port 67. The configuration parameters are carried as options in the DHCP packet in Type, Length, Value (TLV) format; many options are specified in RFC 2132. To limit the number of parameters that servers must provide, hosts specify the parameters that they require, and the server sends only those parameters. Some common options are shown in the following illustration.




Figure 32. DHCP packet Format

The following table lists common DHCP options.

Option	Number and Description
Subnet Mask	Option 1 Specifies the client's subnet mask.
Router	Option 3 Specifies the router IP addresses that may serve as the client's default gateway.
Domain Name Server	Option 6 Specifies the domain name servers (DNSs) that are available to the client.
Domain Name	Option 15 Specifies the domain name that clients should use when resolving hostnames via DNS.
IP Address Lease Time	Option 51 Specifies the amount of time that the client is allowed to use an assigned IP address.
DHCP Message Type	Option 53 <ul style="list-style-type: none"> • 1: DHCPDISCOVER • 2: DHCPOFFER • 3: DHCPREQUEST • 4: DHCPDECLINE • 5: DHCPACK • 6: DHCPNACK • 7: DHCPRELEASE • 8: DHCPINFORM
Parameter Request List	Option 55 Clients use this option to tell the server which parameters it requires. It is a series of octets where each octet is DHCP option code.
Renewal Time	Option 58 Specifies the amount of time after the IP address is granted that the client attempts to renew its lease with the <i>original</i> server.
Rebinding Time	Option 59 Specifies the amount of time after the IP address is granted that the client attempts to renew its lease with <i>any</i> server, if the original server does not respond.
Vendor Class Identifier	Option 60

Option	Number and Description
	Identifiers a user-defined string used by the Relay Agent to forward DHCP client packets to a specific server.
L2 DHCP Snooping	Option 82 Specifies IP addresses for DHCP messages received from the client that are to be monitored to build a DHCP snooping database.
End	Option 255 Signals the last option in the DHCP packet.

Assign an IP Address using DHCP

The following section describes DHCP and the client in a network.

When a client joins a network:

1. The client initially broadcasts a **DHCPDISCOVER** message on the subnet to discover available DHCP servers. This message includes the parameters that the client requires and might include suggested values for those parameters.
2. Servers unicast or broadcast a **DHCPOFFER** message in response to the DHCPDISCOVER that offers to the client values for the requested parameters. Multiple servers might respond to a single DHCPDISCOVER; the client might wait a period of time and then act on the most preferred offer.
3. The client broadcasts a **DHCPREQUEST** message in response to the offer, requesting the offered values.
4. After receiving a DHCPREQUEST, the server binds the clients' unique identifier (the hardware address plus IP address) to the accepted configuration parameters and stores the data in a database called a binding table. The server then broadcasts a **DHCPACK** message, which signals to the client that it may begin using the assigned parameters.
5. When the client leaves the network, or the lease time expires, returns its IP address to the server in a **DHCPRELEASE** message.

There are additional messages that are used in case the DHCP negotiation deviates from the process previously described and shown in the illustration below.

DHCPDECLINE	A client sends this message to the server in response to a DHCPACK if the configuration parameters are unacceptable; for example, if the offered address is already in use. In this case, the client starts the configuration process over by sending a DHCPDISCOVER.
DHCPINFORM	A client uses this message to request configuration parameters when it assigned an IP address manually rather than with DHCP. The server responds by unicast.
DHCPNAK	A server sends this message to the client if it is not able to fulfill a DHCPREQUEST; for example, if the requested address is already in use. In this case, the client starts the configuration process over by sending a DHCPDISCOVER.




Figure 33. Client and Server Messaging

Implementation Information

The following describes DHCP implementation.

- Dell Networking implements DHCP based on RFC 2131 and RFC 3046.
- IP source address validation is a sub-feature of DHCP Snooping; the Dell Networking operating system (OS) uses access control lists (ACLs) internally to implement this feature and as such, you cannot apply ACLs to an interface which has IP source address validation. If you configure IP source address validation on a member port of a virtual local area network (VLAN) and then attempt to apply an access list to the VLAN, the system displays the first line in the following message. If you first apply an ACL to a VLAN and then attempt enable IP source address validation on one of its member ports, the system displays the second line in the following message.

```
% Error: Vlan member has access-list configured.
% Error: Vlan has an access-list configured.
```

NOTE: If you enable DHCP Snooping globally and you have any configured L2 ports, any IP ACL, MAC ACL, or DHCP source address validation ACL does not block DHCP packets.

- The Dell Networking OS provides 40K entries that can be divided between leased addresses and excluded addresses. By extension, the maximum number of pools you can configure depends on the subnet mask that you give to each pool. For example, if all pools were configured for a /24 mask, the total would be 40000/253 (approximately 158). If the subnet is increased, more pools can be configured. The maximum subnet that can be configured for a single pool is /17. The system displays an error message for configurations that exceed the allocated memory.
- The MXL switch supports 4K DHCP Snooping entries.
- All platforms support Dynamic ARP Inspection on 16 VLANs per system. For more information, refer to [Dynamic ARP Inspection](#).

NOTE: If the DHCP server is on the top of rack (ToR) and the VLTi (ICL) is down due to a failed link, when a VLT node is rebooted in BMP (Bare Metal Provisioning) mode, it is not able to reach the DHCP server, resulting in BMP failure.

Configure the System to be a DHCP Server

Configuring the system to be a DHCP server is supported on the MXL switch.

A DHCP server is a network device that has been programmed to provide network configuration parameters to clients upon request. Servers typically serve many clients, making host management much more organized and efficient.

The following table lists the key responsibilities of DHCP servers.

Table 15. DHCP Server Responsibilities

DHCP Server Responsibility	Description
Address Storage and Management	DHCP servers are the owners of the addresses used by DHCP clients. The server stores the addresses and manages their use, keeping track of which addresses have been allocated and which are still available.
Configuration Parameter Storage and Management	DHCP servers also store and maintain other parameters that are sent to clients when requested. These parameters specify in detail how a client is to operate.
Lease Management	DHCP servers use leases to allocate addresses to clients for a limited time. The DHCP server maintains information about each of the leases, including lease length.
Responding To Client Requests	DHCP servers respond to different types of requests from clients, primarily, granting, renewing, and terminating leases.
Providing Administration Services	DHCP servers include functionality that allows an administrator to implement policies that govern how DHCP performs its other tasks.

Configuring the Server for Automatic Address Allocation

Automatic address allocation is an address assignment method by which the DHCP server leases an IP address to a client from a pool of available addresses.

An address pool is a range of IP addresses that the DHCP server may assign. The subnet number indexes the address pools.

To create an address pool, follow these steps.

1. Access the DHCP server CLI context.
CONFIGURATION mode

`ip dhcp server`
2. Create an address pool and give it a name.

DHCP mode

`pool name`

3. Specify the range of IP addresses from which the DHCP server may assign addresses.

DHCP <POOL> mode

`network network/prefix-length`

- `network`: the subnet address.
- `prefix-length`: specifies the number of bits used for the network portion of the address you specify.

The prefix-length range is from 17 to 31.

4. Display the current pool configuration.

DHCP <POOL> mode

`show config`

After an IP address is leased to a client, only that client may release the address. The Dell Networking OS performs a IP + MAC source address validation to ensure that no client can release another clients address. This validation is a default behavior and is separate from IP+MAC source address validation.

Configuration Tasks

To configure DHCP, an administrator must first set up a DHCP server and provide it with configuration parameters and policy information including IP address ranges, lease length specifications, and configuration data that DHCP hosts need.

Configuring the Dell system to be a DHCP server is a three-step process:

1. [Configuring the Server for Automatic Address Allocation](#)
2. [Specifying a Default Gateway](#)
3. [Enabling the DHCP Server](#)

Related Configuration Tasks

- [Configure a Method of Hostname Resolution](#)
- [Creating Manual Binding Entries](#)
- [Debugging the DHCP Server](#)
- [Using DHCP Clear Commands](#)

Excluding Addresses from the Address Pool

The DHCP server assumes that all IP addresses in a DHCP address pool are available for assigning to DHCP clients.

You must specify the IP address that the DHCP server should not assign to clients.

To exclude an address, follow this step.

- Exclude an address range from DHCP assignment. The exclusion applies to all configured pools.

DHCP mode

`excluded-address`

Specifying an Address Lease Time

To specify an address lease time, use the following command.

- Specify an address lease time for the addresses in a pool.
DHCP <POOL>

```
lease {days [hours] [minutes] | infinite}
```

The default is **24 hours**.

Specifying a Default Gateway

The IP address of the default router should be on the same subnet as the client. To specify a default gateway, follow this step.

- Specify default gateway(s) for the clients on the subnet, in order of preference.
DHCP <POOL>

```
default-router address
```

Enabling the DHCP Server

To set up the DHCP Server, you must first enable it. The DHCP server is disabled by default.

1. Enter the DHCP command-line context.
CONFIGURATION mode

```
ip dhcp server
```

2. Enable DHCP server.
DHCP mode

```
no disable
```

The default is **Disabled**.

3. Display the current DHCP configuration.
DHCP mode

```
show config
```

In the following illustration, an IP phone powers Power over Ethernet (PoE) and has acquired an IP address from the Dell Networking system, which is advertising link layer discovery protocol (LLDP)-media endpoint discovery (MED). The leased IP address is displayed using the `show ip dhcp binding` command and confirmed using the `show lldp neighbors` command.




Figure 34. Enabling the DHCP Server

Configure a Method of Hostname Resolution

Dell systems are capable of providing DHCP clients with parameters for two methods of hostname resolution—using DNS or NetBIOS WINS.

Using DNS for Address Resolution

A domain is a group of networks. DHCP clients query DNS IP servers when they need to correlate host names to IP addresses.

1. Create a domain.
DHCP <POOL>

`domain-name name`
2. Specify in order of preference the DNS servers that are available to a DHCP client.
DHCP <POOL>

`dns-server address`

Using NetBIOS WINS for Address Resolution

Windows internet naming service (WINS) is a name resolution service that Microsoft DHCP clients use to correlate host names to IP addresses within a group of networks. Microsoft DHCP clients can be one of four types of NetBIOS nodes: broadcast, peer-to-peer, mixed, or hybrid.

1. Specify the NetBIOS WINS name servers, in order of preference, that are available to Microsoft Dynamic Host Configuration Protocol (DHCP) clients.
DHCP <POOL> mode


`netbios-name-server address`
2. Specify the NetBIOS node type for a Microsoft DHCP client. Dell Networking recommends specifying clients as hybrid.
DHCP <POOL> mode

`netbios-node-type type`

Creating Manual Binding Entries

An address binding is a mapping between the IP address and the media access control (MAC) address of a client.

The DHCP server assigns the client an available IP address automatically, and then creates an entry in the binding table. However, the administrator can manually create an entry for a client; manual bindings are useful when you want to guarantee that a particular network device receives a particular IP address. Manual bindings can be considered single-host address pools. There is no limit on the number of manual bindings, but you can only configure one manual binding per host.

 **NOTE:** The Dell Networking OS does not prevent you from using a network IP as a host IP; be sure to not use a network IP as a host IP.

1. Create an address pool.

DHCP mode

```
pool name
```

2. Specify the client IP address.

```
DHCP <POOL>
```

```
host address
```

3. Specify the client hardware address.

```
DHCP <POOL>
```

```
hardware-address hardware-address type
```

- *hardware-address*: the client MAC address.
- *type*: the protocol of the hardware platform.

The default protocol is **Ethernet**.

Debugging the DHCP Server

To debug the DHCP server, use the following command.

- Display debug information for DHCP server.

EXEC Privilege mode

```
debug ip dhcp server [events | packets]
```

Using DHCP Clear Commands

To clear DHCP binding entries, address conflicts, and server counters, use the following commands.

- Clear DHCP binding entries for the entire binding table.

EXEC Privilege mode.

```
clear ip dhcp binding
```

- Clear a DHCP binding entry for an individual IP address.

EXEC Privilege mode.

```
clear ip dhcp binding ip address
```

- Clear a DHCP address conflict.
EXEC Privilege mode.

```
clear ip dhcp conflict
```

- Clear DHCP server counters.
EXEC Privilege mode.


```
clear ip dhcp server statistics
```

Configure the System to be a Relay Agent

DHCP clients and servers request and offer configuration information via broadcast DHCP messages. Routers do not forward broadcasts, so if there are no DHCP servers on the subnet, the client does not receive a response to its request and therefore cannot access the network.

You can configure an interface on the Dell Networking system to relay the DHCP messages to a specific DHCP server using the `ip helper-address dhcp-address` command from INTERFACE mode, as shown in the following illustration. Specify multiple DHCP servers by using the `ip helper-address dhcp-address` command multiple times.

When you configure the `ip helper-address` command, the system listens for DHCP broadcast messages on port 67. The system rewrites packets received from the client and forwards them via unicast to the DHCP servers; the system rewrites the destination IP address and writes its own address as the relay device. Responses from the server are unicast back to the relay agent on port 67 and the relay agent rewrites the destination address and forwards the packet to the client subnet via broadcast or unicast, depending whether the client has set or cleared the BROADCAST flag in the DHCP Client PDUs.

 **NOTE:** DHCP Relay is not available on Layer 2 interfaces and VLANs.

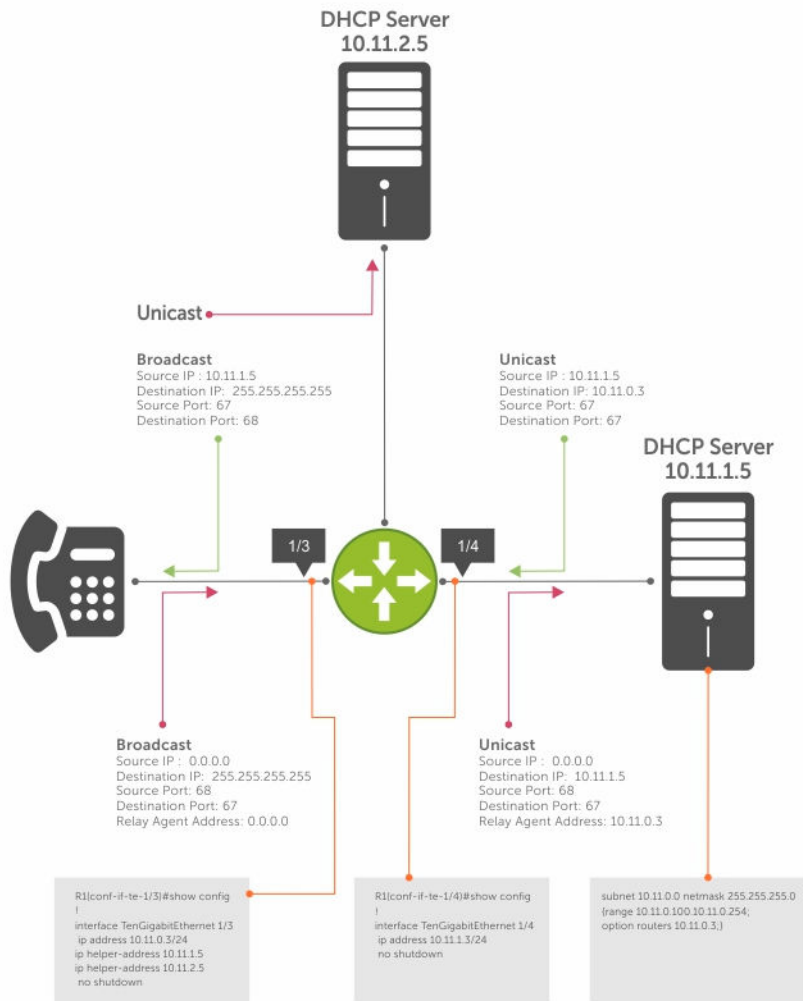


Figure 35. Configuring a Relay Agent

To view the `ip helper-address` configuration for an interface, use the `show ip interface` command from EXEC privilege mode.

Example of the show ip interface Command

```
Dell#show ip int tengig 1/3
GigabitEthernet 1/3 is up, line protocol is down
Internet address is 10.11.0.1/24
Broadcast address is 10.11.0.255
Address determined by user input
IP MTU is 1500 bytes
Helper address is 192.168.0.1
192.168.0.2
Directed broadcast forwarding is disabled
Proxy ARP is enabled
Split Horizon is enabled
Poison Reverse is disabled
```

```
ICMP redirects are not sent
ICMP unreachable are not sent
```

Configure the System to be a DHCP Client

A DHCP client is a network device that requests an IP address and configuration parameters from a DHCP server.

Implement the DHCP client functionality as follows:

- The switch can obtain a dynamically assigned IP address from a DHCP server. A start-up configuration is not received. Use bare metal provisioning (BMP) to receive configuration parameters (the Dell Networking OS version and a configuration file). BMP is enabled as a factory-default setting on a switch.

A switch cannot operate with BMP and as a DHCP client simultaneously. To disable BMP in EXEC mode, use the `stop bmp` command. After BMP stops, the switch acts as a DHCP client.

- By default, the switch is configured to operate in Jumpstart mode as a DHCP client that sends DHCP requests to a DHCP server to retrieve configuration information (IP address, boot-image filename, and configuration file). All ports and management interfaces are brought up in Layer 3 mode and pre-configured with `no shutdown` and `no ip address`. For this reason, you cannot enter configuration commands to set up the switch.

To interrupt a Jumpstart process, prevent a loop from occurring, and apply the FTOS image and startup configuration stored in the local flash, enter the `stop jump-start` command from the console. To re-configure the switch so that it boots up in normal mode using the FTOS image and startup configuration file in local flash, enter the `reload-type normal-reload` command and save it to the startup configuration:

```
FTOS# reload-type normal-reload
FTOS# write memory
FTOS# reload
```

- To re-enable Jumpstart mode for the next reload, enter the `reload-type jump-start` command.
- Acquire a dynamic IP address from a DHCP client is for a limited period or until the client releases the address.
- A DHCP server manages and assigns IP addresses to clients from an address pool stored on the server. For more information, refer to [Configuring the Server for Automatic Address Allocation](#).
- Dynamically assigned IP addresses are supported only on Ethernet interfaces: 10 Gigabit, 40 Gigabit, and 100/1000/10000 Ethernet Interfaces. The DHCP client is supported on VLAN and port-channel interfaces.
- The public out-of-band management interface and default VLAN 1 are configured by default as a DHCP client to acquire a dynamic IP address from a DHCP server.

Configuring the DHCP Client System

This section describes how to configure and view an interface as a DHCP client to receive an IP address.

Dell Networking OS Behavior: The `ip address dhcp` command enables DHCP server-assigned dynamic addresses on an interface. The setting persists after a switch reboot. To stop DHCP transactions and save the dynamically acquired IP address, use the `shutdown` command on the interface. To display the dynamic IP address and show DHCP as the mode of IP address assignment, use the `show interface type slot/port` command. To unconfigure the IP address, use the `no shutdown`

command when the lease timer for the dynamic IP address is expired. The interface acquires a new dynamic IP address from the DHCP server.

If you later enter the `no shutdown` command and the lease timer for the dynamic IP address has expired, the IP address is released.

You cannot configure a secondary (backup) IP address on an interface using the `ip address dhcp` command; you must use the `ip address` command at the interface configuration level.

Use the `no ip address dhcp` command to:


- Release the IP address dynamically acquired from a DHCP server from the interface.
- Disable the DHCP client on the interface so it cannot acquire a dynamic IP address from a DHCP server.
- Stop DHCP packet transactions on the interface.

When you enter the `release dhcp` command, the IP address dynamically acquired from a DHCP server is released from an interface. The ability to acquire a new DHCP server-assigned address remains in the running configuration for the interface. To acquire a new IP address, use the `renew DHCP` command in EXEC Privilege mode or the `ip address dhcp` command in INTERFACE Configuration mode.

To manually configure a static IP address on an interface, use the `ip address` command. A prompt displays to release an existing dynamically acquired IP address. If you confirm, the ability to receive a DHCP server-assigned IP address is removed.

To enable acquiring a dynamic IP address from a DHCP server on an interface configured with a static IP address, use the `ip address dhcp` command. A prompt displays to confirm the IP address reconfiguration. If you confirm, the statically configured IP address is released. An error message displays if you enter the `release dhcp` or `renew dhcp` commands.

To renew the lease time of the dynamically acquired IP, use the `renew dhcp` command on an interface already configured with a dynamic IP address.

 **NOTE:** To verify the currently configured dynamic IP address on an interface, use the `show ip dhcp lease` command. The `show running-configuration` command output only displays **ip address dhcp**. The currently assigned dynamic IP address does not display.

To configure and view an interface as a DHCP client to receive an IP address, use the following commands.

1. Enter INTERFACE Configuration mode on an Ethernet interface.
CONFIGURATION mode

```
interface type slot/port
```

2. Acquire the IP address for an Ethernet interface from a DHCP network server.
INTERFACE mode

```
ip address dhcp
```

Dynamically assigned IP addresses can be released without removing the DHCP client operation on the interface on a switch configured as a DHCP client.

3. Manually acquire a new IP address from the DHCP server by releasing a dynamically acquired IP address while retaining the DHCP client configuration on the interface.

EXEC Privilege mode

```
release dhcp interface type slot/port
```

4. Acquire a new IP address with renewed lease time from a DHCP server.

EXEC Privilege mode

```
renew dhcp interface type slot/port
```

Example of the `show ip dhcp client statistics` Command

Example of the `show ip dhcp lease` command

DHCP Client: Debug Messages Logged during DHCP Client Enabling/Disabling

DHCP Client: Debug Messages Logged during DHCP Client Release/Renew

To display DHCP client information, use the following `show` commands in EXEC Privilege mode.

- To display statistics about DHCP client interfaces, use the `show ip dhcp client statistics interface type slot/port` command.
- To clear DHCP client statistics on a specified or on all interfaces, use the `clear ip dhcp client statistics {all | interface type slot/port}` command.
- To display dynamic IP address lease information currently assigned to a DHCP client interface, use the `show ip dhcp lease [interface type slot/port]` command.
- To display log messages for all DHCP packets sent and received on DHCP client interfaces, use the `debug ip dhcp client packets [interface type slot/port]` command.
- To display log message on DHCP client interfaces for IP address acquisition, IP address release, IP address and lease time renewal, and release an IP address, use the `[no] debug ip dhcp client events [interface type slot/port]` command.

```
Dell# show ip dhcp client statistics interface tengigabitethernet 0/1
```

```
Message      Received
DHCP OFFER   0
DHCPACK      0
DHCPNAK      0
```

```
Message      Sent
DHCPDISCOVER 0
DHCPREQUEST  0
DHCPDECLINE  0
DHCPRELEASE  0
DHCPREBIND   0
DHCPRENEW    0
```

```
Dell# show ip dhcp lease interface tengigabitethernet 4/37
```

```
Interface Lease-IP      Def-Router ServerId  State  Lease Obtnd At  Lease
Expires At
=====
Te 4/37  189.17.9.2/30  0.0.0.0  189.17.9.1  BOUND  06-12-2012 07:35  01-18-2038
11:14
```

```
Renew Time      Rebind Time
=====
09-05-2023 04:56  11-06-2034 13:46
```

The following example shows the packet- and event-level debug messages displayed for the packet transmissions and state transitions on a DHCP client interface when you enable and disable a DHCP client.

```
Dell (conf-if-te-0/1)# ip address dhcp
May 27 15:52:46: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
```



```

Interface Te 0/1 :
DHCP ENABLE CMD Received in state START
May 27 15:52:48: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1:
Transitioned to state SELECTING
May 27 15:52:48: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP DISCOVER sent in Interface Te 0/1
May 27 15:52:48: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
Received DHCP OFFER packet in Interface Te 0/1 with Lease-Ip:10.16.134.250,
Mask:255.255.0.0,Server-Id:10.16.134.249
May 27 15:52:51: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1 :
IP STATUS MESSAGE Received in state SELECTING status: 0
May 27 15:52:51: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1 :
Transitioned to state REQUESTING
May 27 15:52:51: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP REQUEST sent in Interface Te 0/1
May 27 15:52:51: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
Received DHCP ACK packet in Interface Gi 0/1 with Lease-IP:10.16.134.250, Mask:
255.255.0.0,
May 27 15:53:01: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1 :
IP STATUS MESSAGE Received in state REQUESTING status: 0
May 27 15:53:01: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1 :
Transitioned to state BOUND,IP Address: 10.16.134.250 Renewal in 2582 seconds

Dell (conf-if-te-0/1)# no ip address dhcp
May 27 15:53:40: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1 :
DHCP DISABLE CMD Received in state BOUND
May 27 15:53:40: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP RELEASE sent in Interface Te 0/1
May 27 15:53:40: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1 :
Transitioned to state START
May 27 15:53:40: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1 :
DHCP DISABLED CMD sent to Dell in state START

Dell#release dhcp int Te 0/1
Dell#May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT- LOG:
DHCLIENT_DBG_EVT: Interface Te
0/1 :DHCP RELEASE CMD Received in state BOUND
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP RELEASE sent in
Interface Te 0/1
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1
:Transitioned to state STOPPED
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1 :DHCP IP
RELEASED CMD sent to Dell in state STOPPED

Dell#renew dhcp int te 0/1
Dell#May 27 15:55:28: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG:
DHCLIENT_DBG_EVT: Interface Te 0/1 :DHCP
RENEW CMD Received in state STOPPED
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1
:Transitioned to state SELECTING
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP DISCOVER sent in

```

```

Interface Te 0/1
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
Received DHCPOFFER packet
in Interface Te 0/1 with Lease-Ip:10.16.134.250, Mask:255.255.0.0,Server-Id:
10.16.134.249

```

The following shows an example of the packet- and event-level debug messages displayed for the packet transmissions and state transitions on a DHCP client interface when you release and renew a DHCP client.

```

Dell# release dhcp interface tengigabitethernet 0/1
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1 :
DHCP RELEASE CMD Received in state BOUND
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP RELEASE sent in Interface Te 0/1
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1 :
Transitioned to state STOPPED
May 27 15:55:22: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1 :
DHCP IP RELEASED CMD sent to Dell in state STOPPED

```

```

Dell# renew dhcp interface tengigabitethernet 0/1
Dell#May 27 15:55:28: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG:
DHCLIENT_DBG_EVT: Interface Te 0/1 :
DHCP RENEW CMD Received in state STOPPED
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_EVT:
Interface Te 0/1 :
Transitioned to state SELECTING
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
DHCP DISCOVER sent in Interface Te 0/1
May 27 15:55:31: %STKUNIT0-M:CP %DHCLIENT-5-DHCLIENT-LOG: DHCLIENT_DBG_PKT:
Received DHCPOFFER packet in Interface Te 0/1 with Lease-Ip:10.16.134.250,
Mask:255.255.0.0,Server-Id:10.16.134.249


```

DHCP Client on a Management Interface

These conditions apply when you enable a management interface to operate as a DHCP client.

- The management default route is added with the gateway as the router IP address received in the DHCP ACK packet. It is required to send and receive traffic to and from other subnets on the external network. The route is added irrespective when the DHCP client and server are in the same or different subnets. The management default route is deleted if the management IP address is released like other DHCP client management routes.
- *ip route for 0.0.0.0* takes precedence if it is present or added later.
- Management routes added by a DHCP client display with Route Source as **DHCP** in the `show ip management route` and `show ip management-route dynamic` command output.
- Management routes added by DHCP are automatically reinstalled if you configure a static IP route with the `ip route` command that replaces a management route added by the DHCP client. If you remove the statically configured IP route using the `no ip route` command, the management route is reinstalled. Manually delete management routes added by the DHCP client.
- To reinstall management routes added by the DHCP client that is removed or replaced by the same statically configured management routes, release the DHCP IP address and renew it on the management interface.
- Management routes added by the DHCP client have higher precedence over the same statically configured management route. Static routes are not removed from the running configuration if a dynamically acquired management route added by the DHCP client overwrites a static management route.

- Management routes added by the DHCP client are not added to the running configuration.

 **NOTE:** Management routes added by the DHCP client include the specific routes to reach a DHCP server in a different subnet and the management route.

DHCP Client Operation with Other Features

The DHCP client operates with other Dell Networking OS features, as the following describes.

Stacking

The DHCP client daemon runs only on the master unit and handles all DHCP packet transactions. It periodically synchronizes the lease file with the standby unit.

When a stack failover occurs, the new master requires the same DHCP server-assigned IP address on DHCP client interfaces. The new master reinitiates a DHCP packet transaction by sending a DHCP discovery packet on nonbound interfaces.

Virtual Link Trunking (VLT)

A DHCP client is not supported on VLT interfaces.

VLAN and Port Channels

DHCP client configuration and behavior are the same on Virtual LAN (VLAN) and port-channel (LAG) interfaces as on a physical interface.

DHCP Snooping

A DHCP client can run on a switch simultaneously with the DHCP snooping feature as follows:

- If you enable DHCP snooping globally on a switch and you enable a DHCP client on an interface, the trust port, source MAC address, and snooping table validations are not performed on the interface by DHCP snooping for packets destined to the DHCP client daemon.

The following criteria determine packets destined for the DHCP client:

- DHCP is enabled on the interface.
 - The user data protocol (UDP) destination port in the packet is 68.
 - The `chaddr` (change address) in the DHCP header of the packet is the same as the interface's MAC address.
- An entry in the DHCP snooping table is not added for a DHCP client interface.

DHCP Server

A switch can operate as a DHCP client and a DHCP server. DHCP client interfaces cannot acquire a dynamic IP address from the DHCP server running on the switch. Acquire a dynamic IP address from another DHCP server.

Virtual Router Redundancy Protocol (VRRP)

Do not enable the DHCP client on an interface and set the priority to 255 or assign the same DHCP interface IP address to a VRRP virtual group. Doing so guarantees that this router becomes the VRRP group owner.

To use the router as the VRRP owner, if you enable a DHCP client on an interface that is added to a VRRP group, assign a priority less than 255 but higher than any other priority assigned in the group.

Configure Secure DHCP

DHCP as defined by RFC 2131 provides no authentication or security mechanisms.

Secure DHCP is a suite of features that protects networks that use dynamic address allocation from spoofing and attacks.

- [Option 82](#)
- [DHCP Snooping](#)
- [Dynamic ARP Inspection](#)
- [Source Address Validation](#)

Option 82

RFC 3046 (the relay agent information option, or Option 82) is used for class-based IP address assignment.

The code for the relay agent information option is 82, and is comprised of two sub-options, circuit ID and remote ID.

Circuit ID	This is the interface on which the client-originated message is received.
Remote ID	This identifies the host from which the message is received. The value of this sub-option is the MAC address of the relay agent that adds Option 82.

The DHCP relay agent inserts Option 82 before forwarding DHCP packets to the server. The server can use this information to:

- track the number of address requests per relay agent. Restricting the number of addresses available per relay agent can harden a server against address exhaustion attacks.
- associate client MAC addresses with a relay agent to prevent offering an IP address to a client spoofing the same MAC address on a different relay agent.
- assign IP addresses according to the relay agent. This prevents generating DHCP offers in response to requests from an unauthorized relay agent.

The server echoes the option back to the relay agent in its response, and the relay agent can use the information in the option to forward a reply out the interface on which the request was received rather than flooding it on the entire VLAN.

The relay agent strips Option 82 from DHCP responses before forwarding them to the client.

To insert Option 82 into DHCP packets, follow this step.

- Insert Option 82 into DHCP packets.
CONFIGURATION mode

```
ip dhcp relay information-option [trust-downstream]
```

For routers between the relay agent and the DHCP server, enter the `trust-downstream` option.

DHCP Snooping

DHCP snooping protects networks from spoofing. In the context of DHCP snooping, ports are either trusted or not trusted.

By default, all ports are not trusted. Trusted ports are ports through which attackers cannot connect. Manually configure ports connected to legitimate servers and relay agents as trusted.


When you enable DHCP snooping, the relay agent builds a binding table — using DHCPACK messages — containing the client MAC address, IP addresses, IP address lease time, port, VLAN ID, and binding type. Every time the relay agent receives a DHCPACK on a trusted port, it adds an entry to the table.

The relay agent checks all subsequent DHCP client-originated IP traffic (DHCPRELEASE, DHCPNACK, and DHCPDECLINE) against the binding table to ensure that the MAC-IP address pair is legitimate and that the packet arrived on the correct port. Packets that do not pass this check are forwarded to the server for validation. This checkpoint prevents an attacker from spoofing a client and declining or releasing the real client's address. Server-originated packets (DHCP OFFER, DHCPACK, and DHCPNACK) that arrive on a not trusted port are also dropped. This checkpoint prevents an attacker from acting as an imposter as a DHCP server to facilitate a man-in-the-middle attack.

Binding table entries are deleted when a lease expires, or the relay agent encounters a DHCPRELEASE, DHCPNACK, or DHCPDECLINE.

Dell Networking OS Behavior: Introduced in the Dell Networking OS version 7.8.1.0, DHCP snooping was available for Layer 3 only and dependent on DHCP relay agent (`ip helper-address`). The Dell Networking OS version 8.2.1.0 extends DHCP snooping to Layer 2 and you do not have to enable relay agent to snoop on Layer 2 interfaces.

Dell Networking OS Behavior: Binding table entries are deleted when a lease expires or when the relay agent encounters a DHCPRELEASE. The switch maintains a list of snooped VLANs. When the binding table is exhausted, DHCP packets are dropped on snooped VLANs, while these packets are forwarded across non-snooped VLANs. Because DHCP packets are dropped, no new IP address assignments are made. However, DHCPRELEASE and DHCPDECLINE packets are allowed so that the DHCP snooping table can decrease in size. After the table usage falls below the maximum limit of 4000 entries, new IP address assignments are allowed.

 **NOTE:** DHCP server packets are dropped on all not trusted interfaces of a system configured for DHCP snooping. To prevent these packets from being dropped, configure `ip dhcp snooping trust` on the server-connected port.

Enabling DHCP Snooping

To enable DHCP snooping, use the following commands.

1. Enable DHCP snooping globally.
CONFIGURATION mode

```
ip dhcp snooping
```
2. Specify ports connected to DHCP servers as trusted.
INTERFACE mode

```
ip dhcp snooping trust
```

3. Enable DHCP snooping on a VLAN.

CONFIGURATION mode

```
ip dhcp snooping vlan
```

Enabling IPv6 DHCP Snooping

To enable IPv6 DHCP snooping, use the following commands.

1. Enable IPv6 DHCP snooping globally.

CONFIGURATION mode

```
ipv6 dhcp snooping
```

2. Specify ports connected to IPv6 DHCP servers as trusted.

INTERFACE mode

```
ipv6 dhcp snooping trust
```

3. Enable IPv6 DHCP snooping on a VLAN or range of VLANs.

CONFIGURATION mode

```
ipv6 dhcp snooping vlan vlan-id
```

Adding a Static Entry in the Binding Table

To add a static entry in the binding table, use the following command.

- Add a static entry in the binding table.

EXEC Privilege mode

```
ip dhcp snooping binding mac
```

Adding a Static IPV6 DHCP Snooping Binding Table

To add a static entry in the snooping database, use the following command.

- Add a static entry in the snooping binding table.

EXEC Privilege mode

```
ipv6 dhcp snooping binding mac address vlan-id vlan-id ipv6 ipv6-address  
interface interface-type | interface-number lease value
```

Clearing the Binding Table

To clear the binding table, use the following command.

- Delete all of the entries in the binding table.

EXEC Privilege mode

```
clear ip dhcp snooping binding
```

Clearing the DHCP IPv6 Binding Table

To clear the DHCP IPv6 binding table, use the following command.

- Delete all of the entries in the binding table.

EXEC Privilege mode

```
clear ipv6 dhcp snooping binding
```

```
Dell# clear ipv6 dhcp snooping?  
binding Clear the snooping binding database
```

Displaying the Contents of the Binding Table

To display the contents of the binding table, use the following command.

- Display the contents of the binding table.

EXEC Privilege mode

```
show ip dhcp snooping
```

Example of the show ip dhcp snooping Command

View the DHCP snooping statistics with the show ip dhcp snooping command.

```
Dell#show ip dhcp snooping
```

```
IP DHCP Snooping : Disabled.  
IP DHCP Snooping Mac Verification : Disabled.  
IP DHCP Relay Information-option : Disabled.  
IP DHCP Relay Trust Downstream : Enabled.  
  
Database write-delay (In minutes) : 0  
  
DHCP packets information  
Relay Information-option packets : 0  
Relay Trust downstream packets : 0  
Snooping packets : 0  
  
Packets received on snooping disabled L3 Ports : 0  
Snooping packets processed on L2 vlans : 0  
  
DHCP Binding File Details  
Invalid File : 0  
Invalid Binding Entry : 0  
Binding Entry lease expired : 0
```

Displaying the Contents of the DHCPv6 Binding Table

To display the contents of the DHCP IPv6 binding table, use the following command.

- Display the contents of the binding table.

EXEC Privilege mode

```
show ipv6 dhcp snooping biniding
```

Example of the show ipv6 dhcp snooping binding Command

View the DHCP snooping statistics with the show ipv6 dhcp snooping command.

```
Dell#show ipv6 dhcp snooping binding
Codes : S - Static D - Dynamic
IPv6 Address      MAC Address          Expires (Sec)  Type  VLAN    Interface
=====
11:11::22         11:22:11:22:11:22   120331         S     V1 100   Te 0/0
33::22           11:22:11:22:11:23   120331         S     V1 200   Te 0/0
333:22::22       11:22:11:22:11:24   120331         D     V1 300   Te 0/1
```

Debugging the IPv6 DHCP

To debug the IPv6 DHCP, use the following command.

- Display debug information for IPV6 DHCP.
EXEC Privilege mode

```
debug ipv6 dhcp
```

IPv6 DHCP Snooping MAC-Address Verification

Configure to enable verify source mac-address in the DHCP packet against the mac address stored in the snooping binding table.

- Enable IPV6 DHCP snooping .
CONFIGURATION mode

```
ipv6 dhcp snooping verify mac-address
```

Drop DHCP Packets on Snooped VLANs Only

Binding table entries are deleted when a lease expires or the relay agent encounters a DHCPRELEASE.

Starting with the Dell Networking OS version 8.2.1.1, line cards maintain a list of snooped VLANs. When the binding table fills, DHCP packets are dropped only on snooped VLANs, while such packets are forwarded across non-snooped VLANs. Because DHCP packets are dropped, no new IP address assignments are made. However, DHCP release and decline packets are allowed so that the DHCP snooping table can decrease in size. After the table usage falls below the maximum limit of 4000 entries, new IP address assignments are allowed.

To view the number of entries in the table, use the show ip dhcp snooping binding command. This output displays the snooping binding table created using the ACK packets from the trusted port.

```
Dell#show ip dhcp snooping binding
Codes : S - Static D - Dynamic

IP Address  MAC Address          Expires (Sec)  Type  VLAN    Interface
=====
10.1.1.251  00:00:4d:57:f2:50   172800         D     V1 10    Gi 0/2
10.1.1.252  00:00:4d:57:e6:f6   172800         D     V1 10    Gi 0/1
10.1.1.253  00:00:4d:57:f8:e8   172740         D     V1 10    Gi 0/3
10.1.1.254  00:00:4d:69:e8:f2   172740         D     V1 10    Te 0/50
```

```
Total number of Entries in the table : 4
```


Dynamic ARP Inspection

Dynamic address resolution protocol (ARP) inspection prevents ARP spoofing by forwarding only ARP frames that have been validated against the DHCP binding table.


ARP is a stateless protocol that provides no authentication mechanism. Network devices accept ARP requests and replies from any device. ARP replies are accepted even when no request was sent. If a client receives an ARP message for which a relevant entry already exists in its ARP cache, it overwrites the existing entry with the new information.

The lack of authentication in ARP makes it vulnerable to spoofing. ARP spoofing is a technique attackers use to inject false IP-to-MAC mappings into the ARP cache of a network device. It is used to launch man-in-the-middle (MITM), and denial-of-service (DoS) attacks, among others.

A spoofed ARP message is one in which the MAC address in the sender hardware address field and the IP address in the sender protocol field are strategically chosen by the attacker. For example, in an MITM attack, the attacker sends a client an ARP message containing the attacker's MAC address and the gateway's IP address. The client then thinks that the attacker is the gateway, and sends all internet-bound packets to it. Likewise, the attacker sends the gateway an ARP message containing the attacker's MAC address and the client's IP address. The gateway then thinks that the attacker is the client and forwards all packets addressed to the client to it. As a result, the attacker is able to sniff all packets to and from the client.

Other attacks using ARP spoofing include:

- | | |
|--------------------------|---|
| Broadcast | An attacker can broadcast an ARP reply that specifies FF:FF:FF:FF:FF:FF as the gateway's MAC address, resulting in all clients broadcasting all internet-bound packets. |
| MAC flooding | An attacker can send fraudulent ARP messages to the gateway until the ARP cache is exhausted, after which, traffic from the gateway is broadcast. |
| Denial of service | An attacker can send a fraudulent ARP messages to a client to associate a false MAC address with the gateway address, which would blackhole all internet-bound packets from the client. |

 **NOTE:** Dynamic ARP inspection (DAI) uses entries in the L2SysFlow CAM region, a sub-region of SystemFlow. One CAM entry is required for every DAI-enabled VLAN. You can enable DAI on up to 16 VLANs on a system. However, the ExaScale default CAM profile allocates only nine entries to the L2SysFlow region for DAI. You can configure 10 to 16 DAI-enabled VLANs by allocating more CAM space to the L2SysFlow region before enabling DAI.

SystemFlow has 102 entries by default. This region is comprised of two sub-regions: L2Protocol and L2SystemFlow. L2Protocol has 87 entries; L2SystemFlow has 15 entries. Six L2SystemFlow entries are used by Layer 2 protocols, leaving nine for DAI. L2Protocol can have a maximum of 100 entries; you must expand this region to capacity before you can increase the size of L2SystemFlow. This is relevant when you are enabling DAI on VLANs. If, for example, you want to enable DAI on 16 VLANs, you need seven more entries; in this case, reconfigure the SystemFlow region for 122 entries using the `layer-2 eg-acl value fib value frrp value ing-acl value learn value l2pt value qos value system-flow 122` command.

The logic is as follows:

L2Protocol has 87 entries by default and must be expanded to its maximum capacity, 100 entries, before L2SystemFlow can be increased; therefore, 13 more L2Protocol entries are required. L2SystemFlow has 15 entries by default, but only nine are for DAI; to enable DAI on 16 VLANs, seven more entries are required. 87 L2Protocol + 13 additional L2Protocol + 15 L2SystemFlow + 7 additional L2SystemFlow equals 122.

Configuring Dynamic ARP Inspection

To enable dynamic ARP inspection, use the following commands.

1. Enable DHCP snooping.
2. Validate ARP frames against the DHCP snooping binding table.

INTERFACE VLAN mode

```
arp inspection
```

Example of Viewing the ARP Database

Example of Viewing ARP Packets

To view entries in the ARP database, use the `show arp inspection database` command.

```
Dell#show arp inspection database
```

Protocol	Address	Age (min)	Hardware Address	Interface	VLAN	CPU
Internet	10.1.1.251	-	00:00:4d:57:f2:50	Gi 0/2	V1 10	CP
Internet	10.1.1.252	-	00:00:4d:57:e6:f6	Gi 0/1	V1 10	CP
Internet	10.1.1.253	-	00:00:4d:57:f8:e8	Gi 0/3	V1 10	CP
Internet	10.1.1.254	-	00:00:4d:69:e8:f2	Te 0/50	V1 10	CP

Dell#

To see how many valid and invalid ARP packets have been processed, use the `show arp inspection statistics` command.

```
Dell#show arp inspection statistics
```

```
Dynamic ARP Inspection (DAI) Statistics
```

```

-----
Valid ARP Requests      : 0
Valid ARP Replies      : 1000
Invalid ARP Requests   : 1000
Invalid ARP Replies    : 0
Dell#

```

Bypassing the ARP Inspection

You can configure a port to skip ARP inspection by defining the interface as trusted, which is useful in multi-switch environments.

ARPs received on trusted ports bypass validation against the binding table. All ports are untrusted by default.

To bypass the ARP inspection, use the following command.

- Specify an interface as trusted so that ARPs are not validated against the binding table.

```

INTERFACE mode

    arp inspection-trust

```

Dell Networking OS Behavior: Introduced in the Dell Networking OS version 8.2.1.0, DAI was available for Layer 3 only. However, the Dell Networking OS version 8.2.1.1 extends DAI to Layer 2.

Source Address Validation

Using the DHCP binding table, the Dell Networking OS can perform three types of source address validation (SAV).

Table 16. Three Types of Source Address Validation

Source Address Validation	Description
IP Source Address Validation	Prevents IP spoofing by forwarding only IP packets that have been validated against the DHCP binding table.
DHCP MAC Source Address Validation	Verifies a DHCP packet's source hardware address matches the client hardware address field (CHADDR) in the payload.
IP+MAC Source Address Validation	Verifies that the IP source address and MAC source address are a legitimate pair.

Enabling IP Source Address Validation

IP source address validation (SAV) prevents IP spoofing by forwarding only IP packets that have been validated against the DHCP binding table.

A spoofed IP packet is one in which the IP source address is strategically chosen to disguise the attacker. For example, using ARP spoofing, an attacker can assume a legitimate client's identity and receive traffic addressed to it. Then the attacker can spoof the client's IP address to interact with other clients.

The DHCP binding table associates addresses the DHCP servers assign, with the port on which the requesting client is attached. When you enable IP source address validation on a port, the system verifies that the source IP address is one that is associated with the incoming port. If an attacker is impersonating a legitimate client, the source address appears on the wrong ingress port and the system drops the

packet. Likewise, if the IP address is fake, the address is not on the list of permissible addresses for the port and the packet is dropped.

To enable IP source address validation, use the following command.

- Enable IP source address validation.
INTERFACE mode

```
ip dhcp source-address-validation
```

DHCP MAC Source Address Validation

DHCP MAC source address validation (SAV) validates a DHCP packet's source hardware address against the client hardware address field (CHADDR) in the payload. The Dell Networking OS version 8.2.1.1 ensures that the packet's source MAC address is checked against the CHADDR field in the DHCP header only for packets from snooped VLANs.

- Enable DHCP MAC SAV.
CONFIGURATION mode

```
ip dhcp snooping verify mac-address
```

Enabling IP+MAC Source Address Validation

IP source address validation (SAV) validates the IP source address of an incoming packet against the DHCP snooping binding table. IP+MAC SAV ensures that the IP source address and MAC source address are a legitimate pair, rather than validating each attribute individually. You cannot configure IP+MAC SAV with IP SAV.

1. Allocate at least one FP block to the ipmacacl CAM region.
CONFIGURATION mode

```
cam-acl 12acl
```
2. Save the running-config to the startup-config.
EXEC Privilege mode

```
copy running-config startup-config
```
3. Reload the system.
EXEC Privilege

```
reload
```
4. Enable IP+MAC SAV.
INTERFACE mode

```
ip dhcp source-address-validation ipmac
```

The system creates an ACL entry for each IP+MAC address pair in the binding table and applies it to the interface.



To display the IP+MAC ACL for an interface for the entire system, use the `show ip dhcp snooping source-address-validation [interface]` command in EXEC Privilege mode.

Equal Cost Multi-Path (ECMP)

Equal cost multi-path (ECMP) is supported on the MXL switch.

ECMP for Flow-Based Affinity

ECMP for flow-based affinity is available on the MXL switch.

-  **NOTE:** IPv6 /128 routes having multiple paths do not form ECMPs. The /128 route is treated as a host entry and finds its place in the host table.
-  **NOTE:** Using XOR algorithms results in imbalanced loads across an ECMP/LAG when the number of members in said ECMP/LAG is a multiple of 4.


Enabling Deterministic ECMP Next Hop

Deterministic ECMP next hop arranges all ECMPs in order before writing them into the content addressable memory (CAM).

For example, suppose the RTM learns eight ECMPs in the order that the protocols and interfaces came up. In this case, the forwarding information base (FIB) and CAM sort them so that the ECMPs are always arranged. This implementation ensures that every chassis having the same prefixes orders the ECMPs the same.

With eight or less ECMPs, the ordering is lexicographic and deterministic. With more than eight ECMPs, ordering is deterministic, but it is not in lexicographic order.

To enable deterministic ECMP next hop, use the appropriate command.

-  **NOTE:** Packet loss might occur when you enable `ip/ipv6 ecmp-deterministic` for the first-time only.

- Enable IPv4 Deterministic ECMP Next Hop.
CONFIGURATION mode.

```
ip ecmp-deterministic
```

- Enable IPv6 Deterministic ECMP Next Hop.
CONFIGURATION mode.

```
ipv6 ecmp-deterministic
```


Link Bundle Monitoring

Monitoring linked ECMP bundles allows traffic distribution amounts in a link to be monitored for unfair distribution at any given time.

A threshold of 60% is defined as an acceptable amount of traffic on a member link. Links are monitored in 15-second intervals for three consecutive instances. Any deviation within that time causes a syslog to be sent and an alarm event to be generated. When the deviation clears, another syslog is sent and a clear alarm event is generated.

Link bundle utilization is calculated as the total bandwidth of all links divided by the total bytes-per-second of all links. Within each ECMP group, you can specify interfaces. If you enable monitoring for the ECMP group, the utilization calculation is performed when the utilization of the link-bundle (not a link within a bundle) exceeds 60%.

Enable link bundle monitoring using the `ecmp-group` command.

 **NOTE:** An `ecmp-group` index is generated automatically for each unique `ecmp-group` when you configure multipath routes to the same network. The system can generate a maximum of 512 unique `ecmp-groups`. The `ecmp-group` indexes are generated in even numbers (0, 2, 4, 6... 1022) and are for information only.

To enable the link bundle monitoring feature, for link bundle monitoring with ECMP, use the `ecmp-group` command.

You can configure the `ecmp-group` with `id 2`, enabled for link bundle monitoring. This is different from the `ecmp-group` index 2 that is created by configuring routes and is automatically generated. These two `ecmp-groups` are not related in any way.

Managing ECMP Group Paths

Configure the maximum number of paths for an ECMP route that the L3 CAM can hold to avoid path degeneration.

When you do not configure the maximum number of routes, the CAM can hold a maximum ECMP per route.

To configure the maximum number of paths, use the following command.

 **NOTE:** Save the new ECMP settings to the startup-config (`write mem`) then reload the system for the new settings to take effect.

- Configure the maximum number of paths per ECMP group.
CONFIGURATION mode.

```
ip ecmp-group maximum-paths {2-64}
```

- Enable ECMP group path management.
CONFIGURATION mode.

```
ip ecmp-group path-fallback
```

Example of the ip ecmp-group maximum-paths Command

```
Dell(conf)#ip ecmp-group maximum-paths 3
User configuration has been changed. Save the configuration and reload to take
effect
Dell(conf)#
```

FC FLEXIO FPORT

FC FlexIO FPort is now supported on the MXL switch platform.

FC FLEXIO FPORT

The MXL blade switch is a Trident+ based switch which is plugged into the Dell M1000 Blade server chassis. The blade module contains two slots for pluggable flexible module. The goal is to provide support for direct connectivity to FC equipments through Fibre channel ports by FC Flex IO optional module. The FC Flex IO utilizes Broadcom Montreal (BCM84757) FC/FCOE mapper to provide FCOE to FC functionality.

Configuring Switch Mode to FCF Port Mode

To configure switch mode to Fabric services, use the following commands.

1. Configure Switch mode to FCF Port.

CONFIGURATION mode

```
feature fc fport domain id 2
```



NOTE: Enable `remote-fault-signaling rx off` command in FCF FPort mode on interfaces connected to the Compellent and MDF storage devices.

2. Create an FCoE map with the parameters used in the communication between servers and a SAN fabric.

CONFIGURATION mode

```
fcoe-map map-name
```

3. Configure the association between the dedicated VLAN used to carry FCoE traffic between servers and a SAN, and the fabric where the desired storage arrays are installed.

FCOE MAP mode

```
fabric-id fabric-num vlan vlan-id
```

4. Configure the FCoE mapped address prefix (FC-MAP) value which is used to identify FCoE traffic transmitted on the FCoE VLAN for the specified fabric.

FCOE MAP mode

```
fc-map fc-map-value
```

5. Configure the SAN fabric to which the FC port connects by entering the name of the FCoE map applied to the interface.

INTERFACE mode

```
fcoe-map <fcoe-map-name> {tengigabitEthernet slot/port |  
fortygigabitEthernet slot/port}
```

The FCoE map contains FCoE and FC parameter settings (refer to [FCoE Maps](#)). Manually apply the `fcoe-map` to any Ethernet ports used for FCoE.

Name Server

Each participant in the FC environment has a unique ID, which is called the World Wide Name (WWN). This WWN is a 64-bit address.

A Fibre Channel fabric uses another addressing scheme to address the ports in the switched fabric. Each port in the switched fabric is assigned a 24-bit address by the FC switch. When the device logs into the switch fabric on a specific port with its WWN, the switch assigns the port address to that port, and the switch also maintains the correlation between the port address and the WWN address of the device on that port. This function of the switch is implemented by using a name server, a database of objects in which the fabric attached device registers its values.

FC-ID: 24-bit port address. Consists of three parts:

- **Domain:** Address of the switch itself. There are only 239 addresses available for switches in your SAN environment.
- **Area:** Identifies a group of F_Ports. Each group of ports has a different area number, even if there is only one port in the group.
- **Port:** Provides 256 addresses for identifying attached N_Ports and NL_Ports.

The following sequence explains the operation for the attached N_Port:

- N_Port sends a Fabric Login (FLOGI) as it requests a unique 24-bit address from the Fabric Login Server.
- N_Port sends FLOGI to address 0xFFFFFE. Upon success, it obtains a valid address (FCID).
- N_Port sends a Port Login (PLOGI) to inform the Fabric Name Server of its personality and capabilities, this includes WWNN, WWPN.
- N_Port sends PLOGI to address 0xFFFFFC to register this address with the name server.

Command	Description
<code>show fc ns switch</code>	Display all the devices in name server database of the switch.
<code>show fc ns switch brief</code>	Displays the local name server entries — brief version.

The following configurations are applicable only after configuring the switch mode to FCF Port mode using the `feature fc fport domain id 2` command. When you set Switch mode to FCF Port mode, any previously configured `fcoe-map` is removed.


FCoE Maps

To identify the SAN fabric to which FCoE storage traffic is sent, use an FCoE map.

Using an FCoE map, an NPG operates as an FCoE-FC bridge between an FC SAN and FCoE network by providing FCoE-enabled servers and switches with the necessary parameters to log in to a SAN fabric.


An FCoE map applies the following parameters on server-facing Ethernet and fabric-facing FC ports:

- The dedicated FCoE VLAN used to transport FCoE storage traffic.
- The FC-MAP value used to generate a fabric-provided MAC address.
- The association between the FCoE VLAN ID and FC fabric ID where the desired storage arrays are installed. Each Fibre Channel fabric serves as an isolated SAN topology within the same physical network.
- A server uses the priority to select an upstream FCoE forwarder (FCF priority).
- FIP keepalive (FKA) advertisement timeout.

 **NOTE:** In each FCoE map, the fabric ID, FC-MAP value, and FCoE VLAN must be unique. To access one SAN fabric, use one FCoE map.

When you configure an MXL as an NPG, FCoE transit with FIP snooping is automatically enabled and configured using the parameters in the FCoE map applied to server-facing Ethernet and fabric-facing FC interfaces.

After you apply an FCoE map on an FC port, when you enable the port (using the `no shutdown` command), the NPG starts sending FIP multicast advertisements on behalf of the FC port to downstream servers to advertise the availability of a new FCF port on the FCoE VLAN. The FIP advertisement also contains a keepalive message to maintain connectivity between a SAN fabric and downstream servers.

 **NOTE:** After removing and reapplying the fabric map or after modifying the FCoE map, the Fiber Channel (FC) devices do not re-login. To mitigate this issue, you must first run the `shut` command and then the `no shutdown` command on each member interface after you alter the FCOE map.

Creating an FCoE Map

An FCoE map consists of the following elements.

- An association between the dedicated VLAN used to carry FCoE traffic and SAN fabric where the storage arrays are installed. Use a separate FCoE VLAN for each fabric to which FCoE traffic is forwarded. Any non-FCoE traffic sent on a dedicated FCoE VLAN is dropped.
- The FC-MAP value used to generate the fabric-provided MAC address (FPMA). The server uses the FPMA to transmit FCoE traffic to the fabric. You can associate an FC-MAP with only one FCoE VLAN and vice versa.
- FCF priority: a CNA server uses the priority to select an upstream FCoE forwarder (FCF).
- FIP keepalive (FKA) advertisement timeout.

The values for the FCoE VLAN, fabric ID, and FC-MAP must be unique. Apply an FCoE map on downstream server-facing Ethernet ports and upstream fabric-facing Fibre Channel ports.

1. Create an FCoE map which contains parameters used in the communication between servers and a SAN fabric.

CONFIGURATION mode

```
fcoe-map map-name
```

2. Configure the association between the dedicated VLAN and the fabric where the desired storage arrays are installed.

FCoE MAP mode

```
fabric-id fabric-num vlan vlan-id
```

The fabric and VLAN ID numbers must be the same.

The fabric and VLAN ID range is from 2 to 4094.

For example:

```
fabric id 10 vlan 10
```

3. Add a text description of the settings in the FCoE map.

FCoE MAP mode

```
description text
```

The maximum is 32 characters.

4. Specify the FC-MAP value used to generate a fabric-provided MAC address, which is required to send FCoE traffic from a server on the FCoE VLAN to the FC fabric specified in Step 2.

FCoE MAP mode

```
fc-map fc-map-value
```

You must enter a unique MAC address prefix as the FC-MAP value for each fabric.

The range is from 0EFC00 to 0EFCFF.

The default is none.

5. Configure the priority used by a server CNA to select the FCF for a fabric login (FLOGI).

FCoE MAP mode

```
fcf-priority priority
```

The range is from 1 to 255.

The default is **128**.

6. Enable the monitoring FIP keep-alive messages (if it is disabled) to detect if other FCoE devices are reachable.

FCoE MAP mode

```
keepalive
```

The default is FIP keep-alive monitoring is enabled.

7. Configure the time interval (in seconds) used to transmit FIP keepalive advertisements.

FCoE MAP mode

```
fka-adv-period seconds
```

The range is from 8 to 90 seconds.

The default is **8 seconds**.

Zoning

The zoning configurations are supported for Fabric FCF Port mode operation on the MXL.

In FCF Port mode, the `fcoe-map fabric map-name` has the default Zone mode set to deny. This setting denies all the fabric connections unless included in an active zoneset. To change this setting, use the `default-zone-allow` command. Changing this setting to all allows all the fabric connections without zoning.

Zoning is a mechanism to ensure only the nodes that are part of a zone can communicate with each other. Zoning prevents unauthorized access of storage assets. A zone consists of members which are nodes that the adapter address, fabric address, interface, or alias specifies. S5000 supports using World Wide Port Name (WWPN), Fibre Channel ID (FC-ID), or alias as members of a zone.

- WWPN: End device's port WWN name.
- FC-ID: Switch assigned 24-bit device FC address.
- Alias: User-defined name of a zone member.

 **NOTE:** The MXL does not support using WWNN or Domain/Port as members of a zone.

Creating Zone and Adding Members

To create a zone and add members to the zone, use the following commands.

1. Create a zone.

CONFIGURATION mode

```
fc zone zonename
```

2. Add members to a zone.

ZONE CONFIGURATION mode

```
member word
```

The member can be WWPN (00:00:00:00:00:00:00:00), port ID (000000), or alias name (word).

Example of Creating a Zone and Adding Members

```
Dell(conf)#fc zone z1
Dell(conf-fc-zone-z1)#member 11:11:11:11:11:11:11:11
Dell(conf-fc-zone-z1)#member 020202
Dell(conf-fc-zone-z1)#exit
```

Creating Zone Alias and Adding Members

To create a zone alias and add devices to the alias, follow these steps.

1. Create a zone alias name.

CONFIGURATION mode

```
fc alias ZoneAliasName
```

2. Add devices to an alias.

ALIAS CONFIGURATION mode

```
member word
```

The member can be WWPN (00:00:00:00:00:00:00:00), port ID (000000), or alias name (word).

Example of Creating a Zone Alias and Adding Members

```
Dell(conf)#fc alias all
Dell(conf-fc-alias-all)#member 030303
Dell(conf-fc-alias-all)#exit
Dell(conf)#fc zone z1
Dell(conf-fc-zone-z1)#member all
Dell(conf-fc-zone-z1)#exit
```

Creating Zonesets

A zoneset is a grouping or configuration of zones.

To create a zoneset and zones into the zoneset, use the following steps.

1. Create a zoneset.

CONFIGURATION mode

```
fc zoneset zoneset_name
```

2. Add zones into a zoneset.

ZONESET CONFIGURATION mode

```
member zonename
```

Example of Creating Zonesets

```
Dell(conf)#fc zoneset zs1
Dell(conf-fc-zoneset-zs1)#member z1
Dell(conf-fc-zoneset-zs1)#
Dell(conf-fc-zoneset-zs1)#exit
Dell(conf-fc-zoneset-zs1)#
```

Activating a Zoneset

Activating a zoneset makes the zones within it effective.

On a switch, only one zoneset can be active. Any changes in an activated zoneset do not take effect until it is re-activated.

By default, the `fcoe-map fabric map-name` does not have any active zonesets.

1. Enter the `fc-fabric` command in `fcoe-map` to active or de-activate the zoneset.

```
Dell(conf-fcoe-map)#fc-fabric
```

Example:

```
Dell(conf)#fcoe-map map
Dell(conf-fcoe-map)#fc-fabric
Dell(conf-fmap-map-fcfabric)#active-zoneset set
Dell(conf-fmap-map-fcfabric)#no active-zoneset?
active-zoneset
Dell(conf-fmap-map-fcfabric)#no active-zoneset ?
<cr>
Dell(conf-fmap-map-fcfabric)#no active-zoneset
```

2. View the active zoneset.

```
show fc zoneset active
```

Displaying the Fabric Parameters

To display information on switch-wide and interface-specific fabric parameters, use the `show` commands in the following table.

Examples of these `show` commands follow this table.

Command	Description
<code>show config</code>	Displays the fabric parameters.
<code>show fcoe-map</code>	Displays the <code>fcoe-map</code> .
<code>show fc ns switch</code>	Display all the devices in name server database of the switch.
<code>show fc ns switch brief</code>	Display all the devices in name server database of the switch - brief version.
<code>show fc zoneset</code>	Displays the zoneset.
<code>show fc zoneset active</code>	Displays the active zoneset.
<code>show fc zone</code>	Displays the configured zone.
<code>show fc alias</code>	Displays the configured alias.
<code>show fc switch</code>	Displays the FC Switch mode and world wide name.

Example of the show config Command

```
Dell(conf-fcoe-SAN_FABRIC)#show config
!
fcoe-map SAN_FABRIC
  description SAN_FABRIC
  fc-map 0efc00
  fabric-id 1002 vlan 1002
!
fc-fabric
  default-zone-allow all
Dell(conf-fcoe-SAN_FABRIC)#
```

Example of the show fcoe-map Command

```
Dell(conf)#do show fcoe-map
```

```
Fabric Name          map

Fabric Type          Fport
Fabric Id            1002
Vlan Id              1002
Vlan priority        3
FC-MAP               0efc00
FKA-ADV-Period       8
Fcf Priority          128
Config-State         ACTIVE
Oper-State           UP
=====
Switch Config Parameters
=====
DomainID             2
=====
Switch Zoning Parameters
=====
Default Zone Mode:   Deny
Active Zoneset:      set
=====
Members
Fc 0/41
Te 0/29
=====
=====
```

Example of the show fc ns switch Command

```
Dell(conf)#do show fc ns sw
```

```
Total number of devices = 1

Switch Name          28:4e:55:4c:4c:29:00:00
Domain Id            2
Switch Port          4
FC-Id                02:04:03
Port Name            20:01:d4:ae:52:44:37:b2
Node Name            20:00:d4:ae:52:44:37:b2
Class of Service     8
Symbolic Port Name   Broadcom Port0 pWWN 20:01:d4:ae:52:44:37:b2
Symbolic Node Name   Broadcom BCM57810 FCoE 7.6.3.0 7.6.59.0 WIN-
KBFI7FJ2FUH
Port Type            N_Port
```

Example of the show fc ns switch brief Command

```
Dell#show fc ns switch brief
Total number of devices = 1
```

```
Intf# Domain FC-ID      Enode-WWPN              Enode-WWNN
Fc 0/3 1      01:35:00 10:00:8c:7c:ff:17:f8:01 20:00:8c:7c:ff:17:f8:01
Dell#
```

Example of the show fc zoneset Command

```
Dell#show fc zoneset
ZoneSetName ZoneName ZoneMember
=====
fcoe_srv_fc_tgt
      brcd_sanb
                brcd_cna1_wwpn1
                sanb_p2tgt1_wwpn

Active Zoneset: fcoe_srv_fc_tgt
ZoneName ZoneMember
=====
brcd_sanb
      10:00:8c:7c:ff:21:5f:8d
      20:02:00:11:0d:03:00:00
Dell#
```

Example of the show fc zoneset active Command

```
Dell#show fc zoneset active
Active Zoneset: fcoe_srv_fc_tgt
ZoneName ZoneMember
=====
brcd_sanb
      10:00:8c:7c:ff:21:5f:8d
      20:02:00:11:0d:03:00:00
Dell#
```

Example of the show fc zone Command

```
Dell#show fc zone
ZoneName ZoneMember
=====
brcd_sanb
      brcd_cna1_wwpn1
      sanb_p2tgt1_wwpn
Dell#
```

Example of the show fc alias Command

```
Dell(conf)#do show fc alias

ZoneAliasName              ZoneMember
=====
test
                          20:02:d4:ae:52:44:38:4f
                          20:34:78:2b:cb:6f:65:57
```

Example of the show fc switch Command

```
Dell(conf)#do show fc switch
Switch Mode : FPORT
Switch WWN  : 10:00:aa:00:00:00:00:ac
Dell(conf)#
```


FCoE Transit

The Fibre Channel over Ethernet (FCoE) Transit feature is supported on the MXL 10/40GbE switch. When you enable the switch for FCoE transit, the switch functions as a FIP snooping bridge.

 **NOTE:** FCoE transit is not supported on Fibre Channel interfaces.

Fibre Channel over Ethernet

FCoE provides a converged Ethernet network that allows the combination of storage-area network (SAN) and LAN traffic on a Layer 2 link by encapsulating Fibre Channel data into Ethernet frames.

FCoE works with the Ethernet enhancements provided in data center bridging (DCB) to support lossless (no-drop) SAN and LAN traffic. In addition, DCB provides flexible bandwidth sharing for different traffic types, such as LAN and SAN, according to 802.1p priority classes of service. For more information, refer to the [Data Center Bridging \(DCB\)](#) chapter.

Ensure Robustness in a Converged Ethernet Network

Fibre Channel networks used for SAN traffic employ switches that operate as trusted devices. To communicate with other end devices attached to the Fibre Channel network, end devices log into the switch to which they are attached.

Because Fibre Channel links are point-to-point, a Fibre Channel switch controls all storage traffic that an end device sends and receives over the network. As a result, the switch can enforce zoning configurations, ensure that end devices use their assigned addresses, and secure the network from unauthorized access and denial-of-service (DoS) attacks.

To ensure similar Fibre Channel robustness and security with FCoE in an Ethernet cloud network, FIP establishes virtual point-to-point links between FCoE end-devices (server ENodes and target storage devices) and FCoE forwarders (FCFs) over transit FCoE-enabled bridges.

Ethernet bridges commonly provide ACLs that can emulate a point-to-point link by providing the traffic enforcement required to create a Fibre Channel-level of robustness. You can configure ACLs to emulate point-to-point links, providing control over the traffic received or transmitted into the switch. To automatically generate ACLs, use FIP snooping. In addition, FIP serves as a Layer 2 protocol to:

- Operate between FCoE end-devices and FCFs over intermediate Ethernet bridges to prevent unauthorized access to the network and achieve the required security.
- Allow transit Ethernet bridges to efficiently monitor FIP frames passing between FCoE end-devices and an FCF. To dynamically configure ACLs on the bridge to only permit traffic authorized by the FCF, use the FIP snooping data.

FIP enables FCoE devices to discover one another, initialize and maintain virtual links over an Ethernet network, and access storage devices in a storage area network (SAN). FIP satisfies the Fibre Channel requirement for point-to-point connections by creating a unique virtual link for each connection between an FCoE end-device and an FCF via a transit switch.

FIP provides functionality for discovering and logging into an FCF. After discovering and logging in, FIP allows FCoE traffic to be sent and received between FCoE end-devices (ENodes) and the FCF. FIP uses its own EtherType and frame format. The following illustration shows the communication that occurs between an ENode server and an FCoE switch (FCF).

The following table lists the FIP functions.

Table 17. FIP Functions

FIP Function	Description
FIP VLAN discovery	FCoE devices (ENodes) discover the FCoE VLANs on which to transmit and receive FIP and FCoE traffic.
FIP discovery	FCoE end-devices and FCFs are automatically discovered.
Initialization	FCoE devices learn ENodes from the FLOGI and FDISC to allow immediate login and create a virtual link with an FCoE switch.
Maintenance	A valid virtual link between an FCoE device and an FCoE switch is maintained and the LOGO functions properly.

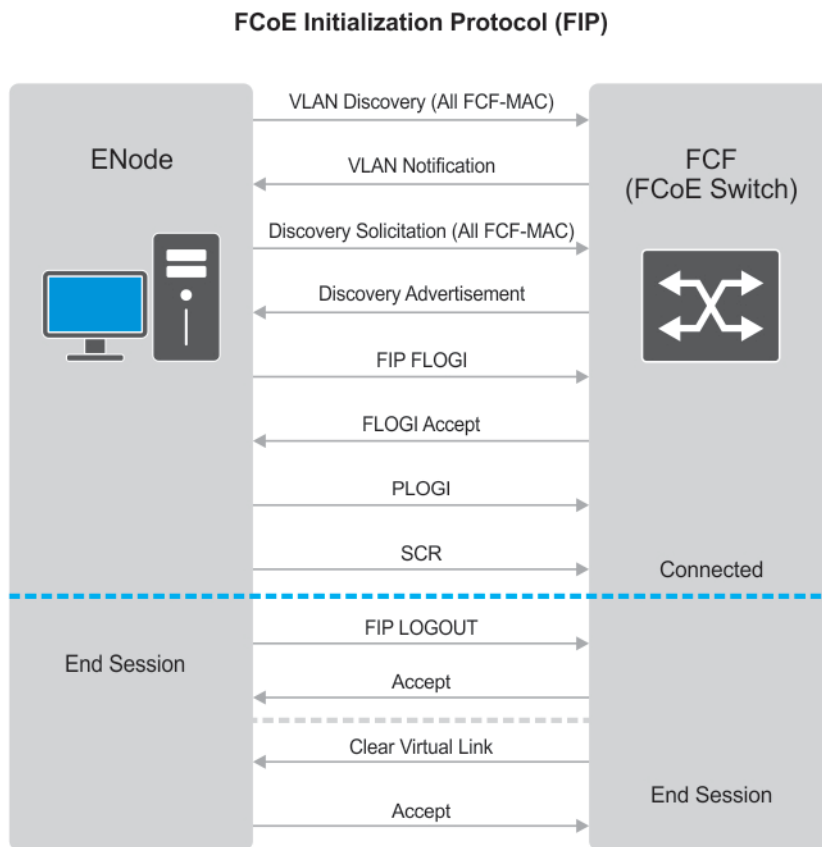


Figure 36. FIP Discovery and Login Between an ENode and an FCF

FIP Snooping on Ethernet Bridges

In a converged Ethernet network, intermediate Ethernet bridges can snoop on FIP packets during the login process on an FCF. Then, using ACLs, a transit bridge can permit only authorized FCoE traffic to be transmitted between an FCoE end-device and an FCF. An Ethernet bridge that provides these functions is called a FIP snooping bridge (FSB).

On a FIP snooping bridge, ACLs are created dynamically as FIP login frames are processed. The ACLs are installed on switch ports configured for ENode mode for server-facing ports and FCF mode for a trusted port directly connected to an FCF.

Enable FIP snooping on the switch and configure the FIP snooping parameters. When you enable FIP snooping, all ports on the switch by default become ENode ports.

Dynamic ACL generation on the MXL switch operating as a FIP snooping bridge function as follows:

Global ACLs These are applied on server-facing ENode ports.

**Port-based
ACLs**

These ACLs are applied on all three port modes: on ports directly connected to an FCF, server-facing ENode ports, and bridge-to-bridge links. Port-based ACLs take precedence over global ACLs.

**FCoE-
generated
ACLs**

These take precedence over user-configured ACLs. A user-configured ACL entry cannot deny FCoE and FIP snooping frames.

The following illustration shows an MXL 10/40GbE switch used as a FIP snooping bridge in a converged Ethernet network. The top-of-rack (ToR) switch operates as an FCF for FCoE traffic. Converged LAN and SAN traffic is transmitted between the ToR switch and an MXL switch. The MXL switch operates as a lossless FIP snooping bridge to transparently forward FCoE frames between the ENode servers and the FCF switch.

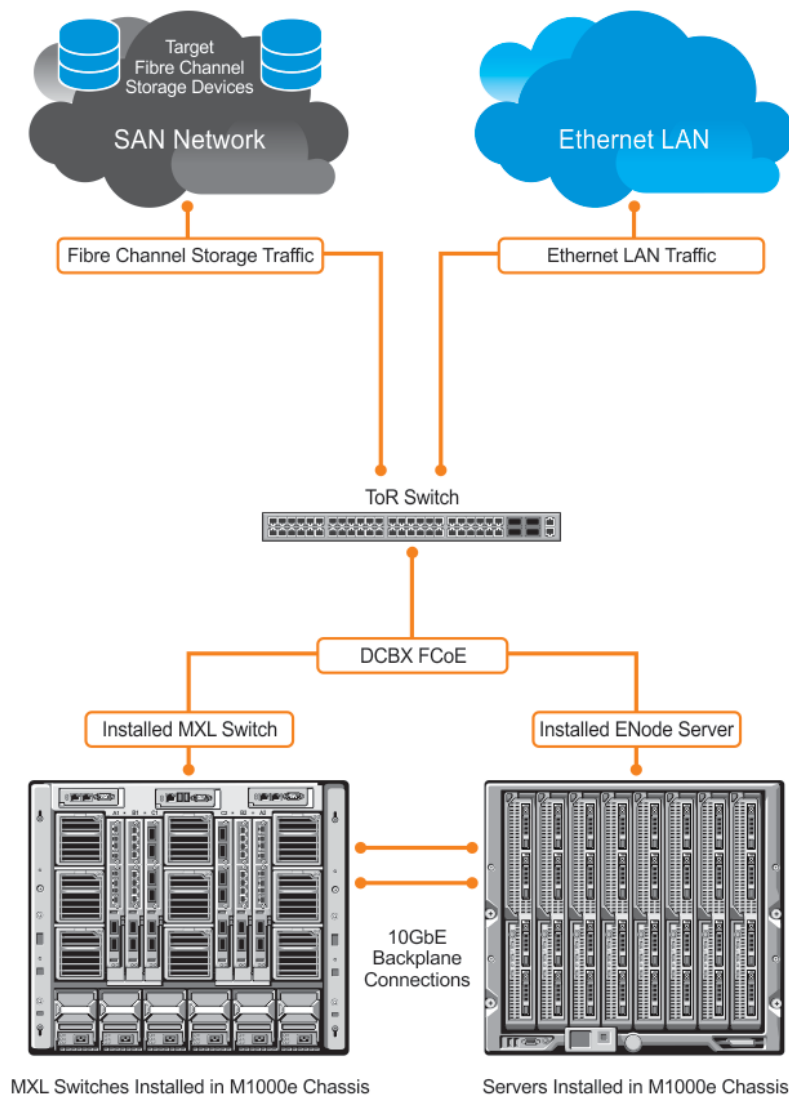


Figure 37. FIP Snooping on an MXL 10/40GbE Switch

The following sections describe how to configure the FIP snooping feature on a switch that functions as a FIP snooping bridge so that it can perform the following functions:

- Perform FIP snooping (allowing and parsing FIP frames) globally on all VLANs or on a per-VLAN basis.
- To assign a MAC address to an FCoE end-device (server ENode or storage device) after a server successfully logs in, set the FCoE MAC address prefix (FC-MAP) value an FCF uses.
- To provide more port security on ports that are directly connected to an FCF and have links to other FIP snooping bridges, set the FCF or Bridge-to-Bridge Port modes.
- To ensure that they are operationally active, check FIP snooping-enabled VLANs.
- Process FIP VLAN discovery requests and responses, advertisements, solicitations, FLOGI/FDISC requests and responses, FLOGO requests and responses, keep-alive packets, and clear virtual-link messages.

FIP Snooping in a Switch Stack

FIP snooping supports switch stacking as follows:

- A switch stack configuration is synchronized with the standby stack unit.
- Dynamic population of the FCoE database (ENode, Session, and FCF tables) is synchronized with the standby stack unit. The FCoE database is maintained by snooping FIP keep-alive messages.
- In case of a failover, the new master switch starts the required timers for the FCoE database tables. Timers run only on the master stack unit.

 **NOTE:** As a best practice, Dell Networking recommends not configuring FIP Snooping on a stacked MXL switch.

Using FIP Snooping

There are four steps to configure FCoE transit.

1. Enable the FCoE transit feature on a switch to maintain FIP snooping information on the switch.
2. Enable FIP snooping globally on all virtual local area networks (VLANs) or individual VLANs on a FIP snooping bridge.
3. Configure the FC-Map value applied globally by the switch on all VLANs or an individual VLAN.
4. Configure FCF mode for a FIP snooping bridge-to-FCF link.


For a sample FIP snooping configuration, refer to [Configuring FIP Snooping](#).

Important Points to Remember

- Enable DCBx on the switch before enabling the FIP Snooping feature.
- To enable the feature on the switch, configure FIP Snooping.
- To allow FIP frames to pass through the switch on all VLANs, enable FIP snooping globally on a switch.
- A switch can support a maximum eight VLANs. Configure at least one FCF/bridge-to-bridge port mode interface for any FIP snooping-enabled VLAN.
- You can configure multiple FCF-trusted interfaces in a VLAN.
- When you disable FIP snooping:
 - ACLs are not installed, FIP and FCoE traffic is not blocked, and FIP packets are not processed.
 - The existing per-VLAN and FIP snooping configuration is stored. The configuration is re-applied the next time you enable the FIP snooping feature.

Enabling the FCoE Transit Feature

The following sections describe how to enable FCoE transit.

 **NOTE:** FCoE transit is disabled by default. To enable this feature, you must follow the [Configuring FIP Snooping](#) procedure.

As soon as you enable the FCoE transit feature on a switch-bridge, existing VLAN-specific and FIP snooping configurations are applied. The FCoE database is populated when the switch connects to a converged network adapter (CNA) or FCF port and compatible DCB configurations are synchronized. By default, all FCoE and FIP frames are dropped unless specifically permitted by existing FIP snooping-generated ACLs. You can reconfigure any of the FIP snooping settings.

If you disable FCoE transit, FIP and FCoE traffic are handled as normal Ethernet frames and no FIP snooping ACLs are generated. The VLAN-specific and FIP snooping configuration is disabled and stored until you re-enable FCoE transit and the configurations are re-applied.

Enable FIP Snooping on VLANs

You can enable FIP snooping globally on a switch on all VLANs or on a specified VLAN.

When you enable FIP snooping on VLANs:

- FIP frames are allowed to pass through the switch on the enabled VLANs and are processed to generate FIP snooping ACLs.
- FCoE traffic is allowed on VLANs only after a successful virtual-link initialization (fabric login FLOGI) between an ENode and an FCF. All other FCoE traffic is dropped.
- You must configure at least one interface for FCF (FIP snooping bridge-bridge) mode on a FIP snooping-enabled VLAN. You can configure multiple FCF trusted interfaces in a VLAN.
- A maximum of eight VLANs are supported for FIP snooping on the switch. When enabled globally, FIP snooping processes FIP packets in traffic only from the first eight incoming VLANs. When enabled on a per-VLAN basis, FIP snooping is supported on up to eight VLANs.

Configure the FC-MAP Value

You can configure the FC-MAP value to be applied globally by the switch on all or individual FCoE VLANs to authorize FCoE traffic.

The configured FC-MAP value is used to check the FC-MAP value for the MAC address assigned to ENodes in incoming FCoE frames. If the FC-MAP value does not match, FCoE frames are dropped. A session between an ENode and an FCF is established by the switch-bridge only when the FC-MAP value on the FCF matches the FC-MAP value on the FIP snooping bridge.

Configure a Port for a Bridge-to-Bridge Link

If a switch port is connected to another FIP snooping bridge, configure the FCoE-Trusted Port mode for bridge-bridge links.

Initially, all FCoE traffic is blocked. Only FIP frames with the ALL_FCF_MAC and ALL_ENODE_MAC values in their headers are allowed to pass. After the switch learns the MAC address of a connected FCF, it allows FIP frames destined to or received from the FCF MAC address.

FCoE traffic is allowed on the port only after the switch learns the FC-MAP value associated with the specified FCF MAC address and verifies that it matches the configured FC-MAP value for the FCoE VLAN.

Configure a Port for a Bridge-to-FCF Link

If a port is directly connected to an FCF, configure the port mode as FCF. Initially, all FCoE traffic is blocked; only FIP frames are allowed to pass.

FCoE traffic is allowed on the port only after a successful fabric login (FLOGI) request/response and confirmed use of the configured FC-MAP value for the VLAN.

FLOGI and fabric discovery (FDISC) request/response packets are trapped to the CPU. They are forwarded after the necessary ACLs are installed.

Impact on Other Software Features

When you enable FIP snooping on a switch, other software features are impacted. The following table lists the impact of FIP snooping.

Table 18. Impact of Enabling FIP Snooping

Impact	Description
MAC address learning	MAC address learning is not performed on FIP and FCoE frames, which are denied by ACLs dynamically created by FIP snooping on server-facing ports in ENode mode.
MTU auto-configuration	MTU size is set to mini-jumbo (2500 bytes) when a port is in Switchport mode, the FIP snooping feature is enabled on the switch, and FIP snooping is enabled on all or individual VLANs.
Link aggregation group (LAG)	FIP snooping is supported on port channels on ports on which PFC mode is on (PFC is operationally up).
STP	If you enable an STP protocol (STP, RSTP, PVSTP, or MSTP) on the switch and ports enter a blocking state, when the state change occurs, the corresponding port-based ACLs are deleted. If a port is enabled for FIP snooping in ENode or FCF mode, the ENode/FCF MAC-based ACLs are deleted.

FIP Snooping Prerequisites

Before you enable FCoE transit and configure FIP snooping on a switch, ensure that certain conditions are met.

A FIP snooping bridge requires data center bridging exchange protocol (DCBx) and priority-based flow control (PFC) to be enabled on the switch for lossless Ethernet connections (refer to the [Data Center](#)

[Bridging \(DCB\)](#) chapter). Dell Networking recommends also enabling enhanced transmission selection (ETS); however, ETS is recommended but not required.

If you enable DCBx and PFC mode is on (PFC is operationally up) in a port configuration, FIP snooping is operational on the port. If the PFC parameters in a DCBx exchange with a peer are not synchronized, FIP and FCoE frames are dropped on the port after you enable the FIP snooping feature.

For VLAN membership, you must:

- create the VLANs on the switch which handles FCoE traffic (use the `interface vlan` command).
- configure each FIP snooping port to operate in Hybrid mode so that it accepts both tagged and untagged VLAN frames (use the `portmode hybrid` command).
- configure tagged VLAN membership on each FIP snooping port that sends and receives FCoE traffic and has links with an FCF, ENode server, or another FIP snooping bridge (use the `tagged port-type slot/port` command).

The default VLAN membership of the port must continue to operate with untagged frames. FIP snooping is not supported on a port that is configured for non-default untagged VLAN membership.

FIP Snooping Restrictions

The following restrictions apply when you configure FIP snooping.

- The maximum number of FCoE VLANs supported on the switch is eight.
- The maximum number of FIP snooping sessions supported per ENode server is 32. To increase the maximum number of sessions to 64, use the `fip-snooping max-sessions-per-enodemac` command.
- The maximum number of FCFs supported per FIP snooping-enabled VLAN is 12.
- Links to other FIP snooping bridges on a FIP snooping-enabled port (bridge-to-bridge links) are not supported on the MXL switch.

Configuring FIP Snooping

You can enable FIP snooping globally on all FCoE VLANs on a switch or on an individual FCoE VLAN. By default, FIP snooping is disabled.

To enable FCoE transit on the switch and configure the FCoE transit parameters on ports, follow these steps.

1. Enable the FCoE transit feature on a switch.

CONFIGURATION mode.

```
feature fip-snooping
```

2. Enable FIP snooping on all VLANs or on a specified VLAN.

CONFIGURATION mode or VLAN INTERFACE mode.

```
fip-snooping enable
```

By default, FIP snooping is disabled on all VLANs.

3. Configure the FC-MAP value used by FIP snooping on all VLANs.

CONFIGURATION VLAN or INTERFACE mode

```
fip-snooping fc-map fc-map-value
```

The default is 0x0EFC00.

The valid values are from 0EFC00 to 0EFCFF.


4. Enter interface configuration mode to configure the port for FIP snooping links.
CONFIGURATION mode

```
interface port-type slot/port
```

By default, a port is configured for bridge-to-ENode links.

5. Configure the port for bridge-to-FCF links.
INTERFACE or CONFIGURATION mode

```
fip-snooping port-mode fcf
```

-  **NOTE:** To disable the FIP snooping feature or FIP snooping on VLANs, use the `no` version of a command; for example, `no feature fip-snooping` or `no fip-snooping enable`.

Displaying FIP Snooping Information

Use the following `show` commands to display information on FIP snooping, .

Table 19. Displaying FIP Snooping Information

Command	Output
<pre>show fip-snooping sessions [interface vlan vlan-id]</pre>	Displays information on FIP-snooped sessions on all VLANs or a specified VLAN, including the ENode interface and MAC address, the FCF interface and MAC address, VLAN ID, FCoE MAC address and FCoE session ID number (FC-ID), worldwide node name (WWNN) and the worldwide port name (WWPN).
<pre>show fip-snooping config</pre>	Displays the FIP snooping status and configured FC-MAP values.
<pre>show fip-snooping enode [enode-mac- address]</pre>	Displays information on the ENodes in FIP-snooped sessions, including the ENode interface and MAC address, FCF MAC address, VLAN ID and FC-ID.
<pre>show fip-snooping fcf [fcf-mac-address]</pre>	Displays information on the FCFs in FIP-snooped sessions, including the FCF interface and MAC address, FCF interface, VLAN ID, FC-MAP value, FKA advertisement period, and number of ENodes connected.
<pre>clear fip-snooping database interface vlan vlan-id {fcoe-mac-address enode- mac-address fcf-mac-address}</pre>	Clears FIP snooping information on a VLAN for a specified FCoE MAC address, ENode MAC address, or FCF MAC address, and removes the corresponding ACLs generated by FIP snooping.

Command	Output
<code>show fip-snooping statistics [interface vlan <i>vlan-id</i> interface <i>port-type port/slot</i> interface <i>port-channel port-channel-number</i>]</code>	Displays statistics on the FIP packets snooped on all interfaces, including VLANs, physical ports, and port channels.
<code>clear fip-snooping statistics [interface vlan <i>vlan-id</i> interface <i>port-type port/slot</i> interface <i>port-channel port-channel-number</i>]</code>	Clears the statistics on the FIP packets snooped on all VLANs, a specified VLAN, or a specified port interface.
<code>show fip-snooping system</code>	Displays information on the status of FIP snooping on the switch (enabled or disabled), including the number of FCoE VLANs, FCFs, ENodes, and currently active sessions.
<code>show fip-snooping vlan</code>	Displays information on the FCoE VLANs on which FIP snooping is enabled.

Example of the `show fip-snooping sessions` Command

Example of Viewing FIP Snooping Configuration

Example of the `show fip-snooping enode` Command

Example of the `show fip-snooping fcf` Command

Example of the `show fip-snooping statistics interface vlan (VLAN and Port)` Command

Example of the `show fip-snooping statistics port-channel` Command

Example of the `show fip-snooping system` Command

Example of the `show fip-snooping vlan` Command

```
Dell#show fip-snooping sessions
Enode MAC           Enode Intf  FCF MAC           FCF Intf  VLAN
aa:bb:cc:00:00:00  Te 0/42    aa:bb:cd:00:00:00 Te 0/43    100
aa:bb:cc:00:00:00  Te 0/42    aa:bb:cd:00:00:00 Te 0/43    100
aa:bb:cc:00:00:00  Te 0/42    aa:bb:cd:00:00:00 Te 0/43    100
aa:bb:cc:00:00:00  Te 0/42    aa:bb:cd:00:00:00 Te 0/43    100
aa:bb:cc:00:00:00  Te 0/42    aa:bb:cd:00:00:00 Te 0/43    100

FCoE MAC           FC-ID       Port WWPN          Port WWNN
0e:fc:00:01:00:01  01:00:01   31:00:0e:fc:00:00:00:00  21:00:0e:fc:00:00:00:00
0e:fc:00:01:00:02  01:00:02   41:00:0e:fc:00:00:00:00  21:00:0e:fc:00:00:00:00
0e:fc:00:01:00:03  01:00:03   41:00:0e:fc:00:00:00:01  21:00:0e:fc:00:00:00:00
0e:fc:00:01:00:04  01:00:04   41:00:0e:fc:00:00:00:02  21:00:0e:fc:00:00:00:00
0e:fc:00:01:00:05  01:00:05   41:00:0e:fc:00:00:00:03  21:00:0e:fc:00:00:00:00
```

The following table describes the `show fip-snooping sessions` command fields.

Table 20. `show fip-snooping sessions` Command Description

Field	Description
ENode MAC	MAC address of the ENode .
ENode Interface	Slot/ port number of the interface connected to the ENode.

Field	Description
FCF MAC	MAC address of the FCF.
FCF Interface	Slot/ port number of the interface to which the FCF is connected.
VLAN	VLAN ID number used by the session.
FCoE MAC	MAC address of the FCoE session assigned by the FCF.
FC-ID	Fibre Channel ID assigned by the FCF.
Port WWPN	Worldwide port name of the CNA port.
Port WWNN	Worldwide node name of the CNA port.

```
Dell# show fip-snooping config
FIP Snooping Feature enabled Status: Enabled
FIP Snooping Global enabled Status: Enabled
Global FC-MAP Value: 0X0EFC00
```

```
FIP Snooping enabled VLANs
VLAN   Enabled  FC-MAP
-----
100    TRUE      0X0EFC00
```

```
Dell# show fip-snooping enode
Enode MAC           Enode Interface FCF MAC           VLAN  FC-ID
-----
d4:ae:52:1b:e3:cd  Te 0/11         54:7f:ee:37:34:40 100   62:00:11
```

The following table describes the show fip-snooping enode command fields.

Table 21. show fip-snooping enode Command Description

Field	Description
ENode MAC	MAC address of the ENode.
ENode Interface	Slot/ port number of the interface connected to the ENode.
FCF MAC	MAC address of the FCF.
VLAN	VLAN ID number used by the session.
FC-ID	Fibre Channel session ID assigned by the FCF.

```
Dell# show fip-snooping fcf
FCF MAC           FCF Interface VLAN  FC-MAP   FKA_ADV_PERIOD  No. of Enodes
-----
54:7f:ee:37:34:40 Po 22         100  0e:fc:00 4000          2
```

The following table describes the show fip-snooping fcf command fields.

Table 22. show fip-snooping fcf Command Description

Field	Description
FCF MAC	MAC address of the FCF.
FCF Interface	Slot/port number of the interface to which the FCF is connected.
VLAN	VLAN ID number used by the session.
FC-MAP	FC-Map value advertised by the FCF.
ENode Interface	Slot/number of the interface connected to the ENode.
FKA_ADV_PERIOD	Period of time (in milliseconds) during which FIP keep-alive advertisements are transmitted.
No of ENodes	Number of ENodes connected to the FCF.
FC-ID	Fibre Channel session ID assigned by the FCF.

```
Dell# show fip-snooping statistics interface vlan 100
Number of Vlan Requests           :0
Number of Vlan Notifications      :0
Number of Multicast Discovery Solicits :2
Number of Unicast Discovery Solicits :0
Number of FLOGI                   :2
Number of FDISC                    :16
Number of FLOGO                    :0
Number of Enode Keep Alive         :9021
Number of VN Port Keep Alive       :3349
Number of Multicast Discovery Advertisement :4437
Number of Unicast Discovery Advertisement :2
Number of FLOGI Accepts            :2
Number of FLOGI Rejects            :0
Number of FDISC Accepts            :16
Number of FDISC Rejects            :0
Number of FLOGO Accepts            :0
Number of FLOGO Rejects            :0
Number of CVL                      :0
Number of FCF Discovery Timeouts    :0
Number of VN Port Session Timeouts :0
Number of Session failures due to Hardware Config :0
Dell(conf)#
```

```
Dell# show fip-snooping statistics int tengigabitethernet 0/11
Number of Vlan Requests           :1
Number of Vlan Notifications      :0
Number of Multicast Discovery Solicits :1
Number of Unicast Discovery Solicits :0
Number of FLOGI                   :1
Number of FDISC                    :16
Number of FLOGO                    :0
Number of Enode Keep Alive         :4416
Number of VN Port Keep Alive       :3136
Number of Multicast Discovery Advertisement :0
Number of Unicast Discovery Advertisement :0
Number of FLOGI Accepts            :0
Number of FLOGI Rejects            :0
Number of FDISC Accepts            :0
Number of FDISC Rejects            :0
```

```

Number of FLOGO Accepts :0
Number of FLOGO Rejects :0
Number of CVL :0
Number of FCF Discovery Timeouts :0
Number of VN Port Session Timeouts :0
Number of Session failures due to Hardware Config :0

```

```

Dell# show fip-snooping statistics interface port-channel 22
Number of Vlan Requests :0
Number of Vlan Notifications :2
Number of Multicast Discovery Solicits :0
Number of Unicast Discovery Solicits :0
Number of FLOGI :0
Number of FDISC :0
Number of FLOGO :0
Number of Enode Keep Alive :0
Number of VN Port Keep Alive :0
Number of Multicast Discovery Advertisement :4451
Number of Unicast Discovery Advertisement :2
Number of FLOGI Accepts :2
Number of FLOGI Rejects :0
Number of FDISC Accepts :16
Number of FDISC Rejects :0
Number of FLOGO Accepts :0
Number of FLOGO Rejects :0
Number of CVL :0
Number of FCF Discovery Timeouts :0
Number of VN Port Session Timeouts :0
Number of Session failures due to Hardware Config :0

```


The following table describes the show fip-snooping statistics command fields.

Table 23. show fip-snooping statistics Command Descriptions

Field	Description
Number of VLAN Requests	Number of FIP-snooped VLAN request frames received on the interface.
Number of VLAN Notifications	Number of FIP-snooped VLAN notification frames received on the interface.
Number of Multicast Discovery Solicits	Number of FIP-snooped multicast discovery solicit frames received on the interface.
Number of Unicast Discovery Solicits	Number of FIP-snooped unicast discovery solicit frames received on the interface.
Number of FLOGI	Number of FIP-snooped FLOGI request frames received on the interface.
Number of FDISC	Number of FIP-snooped FDISC request frames received on the interface.
Number of FLOGO	Number of FIP-snooped FLOGO frames received on the interface.
Number of ENode Keep Alives	Number of FIP-snooped ENode keep-alive frames received on the interface.

Field	Description
Number of VN Port Keep Alives	Number of FIP-snooped VN port keep-alive frames received on the interface.
Number of Multicast Discovery Advertisements	Number of FIP-snooped multicast discovery advertisements received on the interface.
Number of Unicast Discovery Advertisements	Number of FIP-snooped unicast discovery advertisements received on the interface.
Number of FLOGI Accepts	Number of FIP FLOGI accept frames received on the interface.
Number of FLOGI Rejects	Number of FIP FLOGI reject frames received on the interface.
Number of FDISC Accepts	Number of FIP FDISC accept frames received on the interface.
Number of FDISC Rejects	Number of FIP FDISC reject frames received on the interface.
Number of FLOGO Accepts	Number of FIP FLOGO accept frames received on the interface.
Number of FLOGO Rejects	Number of FIP FLOGO reject frames received on the interface.
Number of CVLs	Number of FIP clear virtual link frames received on the interface.
Number of FCF Discovery Timeouts	Number of FCF discovery timeouts that occurred on the interface.
Number of VN Port Session Timeouts	Number of VN port session timeouts that occurred on the interface.
Number of Session failures due to Hardware Config	Number of session failures due to hardware configuration that occurred on the interface.

```
Dell# show fip-snooping system
Global Mode           : Enabled
FCOE VLAN List (Operational) : 1, 100
FCFs                  : 1
Enodes                : 2
Sessions              : 17
```

 **NOTE:** NPIV sessions are included in the number of FIP-snooped sessions displayed.

```
Dell# show fip-snooping vlan
* = Default VLAN

VLAN  FC-MAP    FCFs  Enodes  Sessions
----  -
*1    -          -     -       -
100   0X0EFC00  1     2       17
```

FCoE Transit Configuration Example

The following illustration shows an MXL switch used as a FIP snooping bridge for FCoE traffic between an ENode (server blade) and an FCF (ToR switch). The ToR switch operates as an FCF and FCoE gateway.

In this example, DCBx and PFC are enabled on the FIP snooping bridge and on the FCF ToR switch. On the FIP snooping bridge, DCBx is configured as follows:

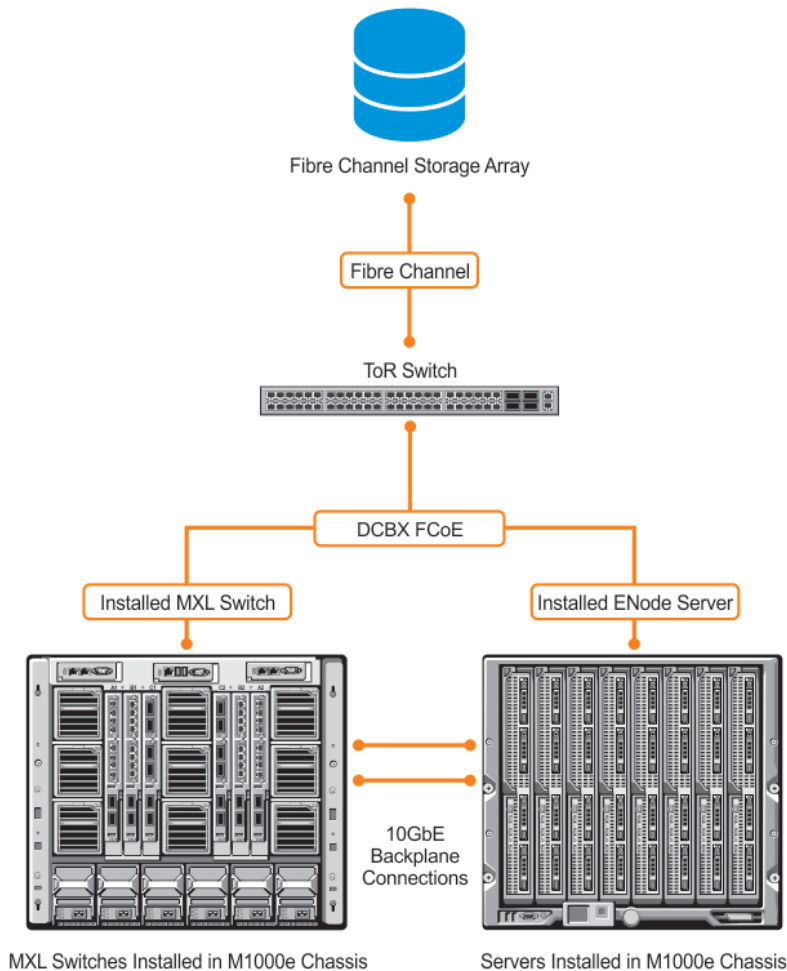


Figure 38. FIP Snooping on an MXL 10/40GbE Switch Configuration Example

- A server-facing port is configured for DCBx in an auto-downstream role.
- An FCF-facing port is configured for DCBx in an auto-upstream or configuration-source role.

The DCBx configuration on the FCF-facing port is detected by the server-facing port and the DCB PFC configuration on both ports is synchronized. For more information about how to configure DCBx and PFC on a port, refer to the [Data Center Bridging \(DCB\)](#) chapter.

The following example shows how to configure FIP snooping on FCoE VLAN 10, on an FCF-facing port (0/50), on an ENode server-facing port (0/1), and to configure the FIP snooping ports as tagged members of the FCoE VLAN enabled for FIP snooping.

Example of Enabling the FIP Snooping Feature on the Switch (FIP Snooping Bridge)

Example of Enabling FIP Snooping on the FCoE VLAN

Example of Enabling an FC-MAP Value on a VLAN

Example of Configuring the ENode Server-Facing Port


Example of Configuring the FCF-Facing Port

Example of Configuring FIP Snooping Ports as Tagged Members of the FCoE VLAN


```
Dell(conf)# feature fip-snooping
```

```
Dell(conf)# interface vlan 10
Dell(conf-if-vl-10)# fip-snooping enable
```

```
Dell(conf-if-vl-10)# fip-snooping fc-map 0x0EFC01
```

 **NOTE:** Configuring an FC-MAP value is only required if you do not use the default FC-MAP value (0x0EFC00).

```
Dell(conf)# interface tengigabitethernet 0/1
Dell(conf-if-te-0/1)# portmode hybrid
Dell(conf-if-te-0/1)# switchport
```

 **NOTE:** A port is enabled by default for bridge-ENode links.


```
Dell(conf)# interface tengigabitethernet 0/50
Dell(conf-if-te-0/50)# portmode hybrid
Dell(conf-if-te-0/50)# switchport
Dell(conf-if-te-0/50)# fip-snooping port-mode fcf
```


```
Dell(conf)# interface vlan 10
Dell(conf-if-vl-10)# tagged tengigabitethernet 0/1
Dell(conf-if-vl-10)# tagged tengigabitethernet 0/50
Dell(conf-if-te-0/1)# no shut
Dell(conf-if-te-0/50)# no shut
Dell(conf-if-vl-10)# no shut
```

After FIP packets are exchanged between the ENode and the switch, a FIP snooping session is established. ACLs are dynamically generated for FIP snooping on the FIP snooping bridge/switch.

FIPS Cryptography

Federal information processing standard (FIPS) cryptography is supported on the MXL switch platform. This chapter describes how to enable FIPS cryptography requirements on Dell Networking platforms. This feature provides cryptographic algorithms conforming to various FIPS standards published by the National Institute of Standards and Technology (NIST), a non-regulatory agency of the US Department of Commerce. FIPS mode is also validated for numerous platforms to meet the FIPS-140-2 standard for a software-based cryptographic module.

 **NOTE:** The FIPS mode included in this release is the OpenSSL FIPS Object Module v2.0, which has been validated to meet FIPS-140-2 requirements, per certificate #1747. The MXL switch platform is not one of the validated platforms. Dell Networking has contracted with the OpenSSL Foundation to complete a Change Letter validation of the MXL switch platform for this FIPS mode. A patch release will be available after that Change Letter validation has been completed.

 **NOTE:** For the Dell Networking OS version 8.3.12.0, only the SSH and SCP copy features use FIPS Cryptographic mode to secure management interface user sessions and file transfers. Other features that use cryptographic algorithms do not, or cannot, use FIPS mode. You must configure the management interfaces to limit access to/from the system to SSH alone.

Preparing the System

Before you enable FIPS mode, Dell Networking recommends making the following changes to your system.

1. Disable the Telnet server (only use secure shell [SSH] to access the system).
2. Disable the FTP server (only use secure copy [SCP] to transfer files to and from the system).
3. Attach a secure, standalone host to the console port for the FIPS configuration to use.

Enabling FIPS Mode


To enable or disable FIPS mode, use the console port.

Secure the host attached to the console port against unauthorized access. Any attempts to enable or disable FIPS mode from a virtual terminal session are denied.

When you enable FIPS mode, the following actions are taken:

- If enabled, the SSH server is disabled.
- All open SSH and Telnet sessions, as well as all SCP and FTP file transfers, are closed.
- Any existing host keys (both RSA and RSA1) are deleted from system memory and NVRAM storage.

- FIPS mode is enabled.
 - If you enable the SSH server when you enter the `fips mode enable` command, it is re-enabled for version 2 *only*.
 - If you re-enable the SSH server, a new RSA host key-pair is generated automatically. You can also manually create this key-pair using the `crypto key generate` command.

 **NOTE:** Under certain unusual circumstances, it is possible for the `fips enable` command to indicate a failure.

- This failure occurs if any of the self-tests fail when you enable FIPS mode.
- This failure occurs if there were existing SSH/Telnet sessions that could not be closed successfully in a reasonable amount of time. In general, this failure can occur if a user at a remote host is in the process of establishing an SSH session to the local system, and has been prompted to accept a new host key or to enter a password, but is not responding to the request. Assuming this failure is a transient condition, attempting to enable FIPS mode again should be successful.

To enable FIPS mode, use the following command.

- Enable FIPS mode from a console port.

```
CONFIGURATION
```

```
fips mode enable
```

The following warning message displays:

```
WARNING: Enabling FIPS mode will close all SSH/Telnet connections, restart
those servers, and destroy all configured host keys. Proceed (y/n) ?
```

Generating Host-Keys

The following describes hot-key generation.

When you enable or disable FIPS mode, the system deletes the current public/private host-key pair, terminates any SSH sessions that are in progress (deleting all the per-session encryption key information), actually enables/tests FIPS mode, generates new host-keys, and re-enables the SSH server (assuming it was enabled before enabling FIPS).

For more information, refer to the *SSH Server and SCP Commands* section in the *Security* chapter of the *Dell Networking OS Command Line Reference Guide*.

Monitoring FIPS Mode Status

To view the status of the current FIPS mode (enabled/disabled), use the following commands.

- Use either command to view the status of the current FIPS mode.

```
show fips status
```

```
show system
```

Example of the show fips status Command

Example of the show system Command

```
Dell#show fips status
FIPS Mode : Enabled
for the system using the show system command.
```

```
Dell#show system

Stack MAC : 00:01:e8:8a:ff:0c

Reload Type : normal-reload [Next boot : normal-reload]

-- Unit 0 --
Unit Type      : Management Unit
Status        : online
Next Boot     : online
Required Type  : XML - 52-port GE/TE/FG (SE)
Current Type  : XML - 52-port GE/TE/FG (SE)
Master priority : 0
Hardware Rev   : 3.0
Num Ports     : 64
Up Time       : 7 hr, 3 min
Dell Version   : XML-8-3-7-1061
Jumbo Capable : yes
POE Capable   : no
FIPS Mode     : enabled
Burned In MAC : 00:01:e8:8a:ff:0c
No Of MACs    : 3
...
```

Disabling FIPS Mode

The following describes disabling FIPS mode.

When you disable FIPS mode, the following changes occur:

- The SSH server disables.
- All open SSH and Telnet sessions, as well as all SCP and FTP file transfers, close.
- Any existing host keys (both RSA and RSA1) are deleted from system memory and NVRAM storage.
- FIPS mode disables.
- The SSH server re-enables.
- The Telnet server re-enables (if it is present in the configuration).
- New 1024-bit RSA and RSA1 host key-pairs are created.

To disable FIPS mode, use the following command.

- To disable FIPS mode from a console port.
CONFIGURATION mode

```
no fips mode enable
```

The following Warning message displays:

```
WARNING: Disabling FIPS mode will close all SSH/Telnet connections, restart
those servers, and destroy
```

all configured host keys.
Proceed (y/n) ?

Force10 Resilient Ring Protocol (FRRP)

FRRP provides fast network convergence to Layer 2 switches interconnected in a ring topology, such as a metropolitan area network (MAN) or large campuses.

FRRP is similar to what can be achieved with the spanning tree protocol (STP), though even with optimizations, STP can take up to 50 seconds to converge (depending on the size of network and node of failure) may require 4 to 5 seconds to reconverge. FRRP can converge within 150ms to 1500ms when a link in the ring breaks (depending on network configuration).

To operate a deterministic network, a network administrator must run a protocol that converges independently of the network size or node of failure. FRRP is a proprietary protocol that provides this flexibility, while preventing Layer 2 loops. FRRP provides sub-second ring-failure detection and convergence/re-convergence in a Layer 2 network while eliminating the need for running spanning-tree protocol. With its two-way path to destination configuration, FRRP provides protection against any single link/switch failure and thus provides for greater network uptime.

Protocol Overview

FRRP is built on a ring topology.

You can configure up to 255 rings on a system. FRRP uses one Master node and multiple Transit nodes in each ring. There is no limit to the number of nodes on a ring. The Master node is responsible for the intelligence of the Ring and monitors the status of the Ring. The Master node checks the status of the Ring by sending ring health frames (RHF) around the Ring from its Primary port and returning on its Secondary port. If the Master node misses three consecutive RHFs, the Master node determines the ring to be in a failed state. The Master then sends a Topology Change RHF to the Transit Nodes informing them that the ring has changed. This causes the Transit Nodes to flush their forwarding tables, and re-converge to the new network structure.

One port of the Master node is designated the Primary port (P) to the ring; another port is designated as the Secondary port (S) to the ring. In normal operation, the Master node blocks the Secondary port for all non-control traffic belonging to this FRRP group, thereby avoiding a loop in the ring, like STP. Layer 2 switching and learning mechanisms operate per existing standards on this ring.

Each Transit node is also configured with a Primary port and a Secondary port on the ring, but the port distinction is ignored as long as the node is configured as a Transit node. If the ring is complete, the Master node logically blocks all data traffic in the transmit and receive directions on the Secondary port to prevent a loop. If the Master node detects a break in the ring, it unblocks its Secondary port and allows data traffic to be transmitted and received through it. Refer to the following illustration for a simple example of this FRRP topology. Note that ring direction is determined by the Master node's Primary and Secondary ports.




Figure 39. Normal Operating FRRP Topology

A virtual LAN (VLAN) is configured on all node ports in the ring. All ring ports must be members of the Member VLAN and the Control VLAN.

The Member VLAN is the VLAN used to transmit data as described earlier.

The Control VLAN is used to perform the health checks on the ring. The Control VLAN can always pass through all ports in the ring, including the secondary port of the Master node.

Ring Status

The ring failure notification and the ring status checks provide two ways to ensure the ring remains up and active in the event of a switch or port failure.

Ring Checking

At specified intervals, the Master node sends a ring health frame (RHF) through the ring. If the ring is complete, the frame is received on its secondary port and the Master node resets its fail-period timer and continues normal operation.

If the Master node does not receive the RHF before the fail-period timer expires (a configurable timer), the Master node moves from the Normal state to the Ring-Fault state and unblocks its Secondary port. The Master node also clears its forwarding table and sends a control frame to all other nodes, instructing them to also clear their forwarding tables. Immediately after clearing its forwarding table, each node starts learning the new topology.

Ring Failure

If a Transit node detects a link down on any of its ports on the FRRP ring, it immediately sends a link-down control frame on the Control VLAN to the Master node.

When the Master node receives this control frame, the Master node moves from the Normal state to the Ring-Fault state and unblocks its Secondary port. The Master node clears its routing table and sends a control frame to all other ring nodes, instructing them to clear their routing tables as well. Immediately after clearing its routing table, each node begins learning the new topology.

Ring Restoration

The Master node continues sending ring health frames out its primary port even when operating in the Ring-Fault state.

After the ring is restored, the next status check frame is received on the Master node's Secondary port. This causes the Master node to transition back to the Normal state. The Master node then logically blocks non-control frames on the Secondary port, clears its own forwarding table, and sends a control frame to the Transit nodes, instructing them to clear their forwarding tables and re-learn the topology.

During the time between the Transit node detecting that its link is restored and the Master node detecting that the ring is restored, the Master node's Secondary port is still forwarding traffic. This can create a temporary loop in the topology. To prevent this, the Transit node places all the ring ports transiting the newly restored port into a temporary blocked state. The Transit node remembers which port has been temporarily blocked and places it into a pre-forwarding state. When the Transit node in the pre-forwarding state receives the control frame instructing it to clear its routing table, it does so and unblocks the previously blocked ring ports on the newly restored port. Then the Transit node returns to the Normal state.

Multiple FRRP Rings

Up to 255 rings are allowed per system and multiple rings can be run on one system.

More than the recommended number of rings may cause interface instability. You can configure multiple rings with a single switch connection; a single ring can have multiple FRRP groups; multiple rings can be connected with a common link.

Member VLAN Spanning Two Rings Connected by One Switch

A member VLAN can span two rings interconnected by a common switch, in a figure-eight style topology.

A switch can act as a Master node for one FRRP group and a Transit for another FRRP group, or it can be a Transit node for both rings.

In the following example, FRRP 101 is a ring with its own Control VLAN, and FRRP 202 has its own Control VLAN running on another ring. A Member VLAN that spans both rings is added as a Member VLAN to both FRRP groups. Switch R3 has two instances of FRRP running on it: one for each ring. The example topology that follows shows R3 assuming the role of a Transit node for both FRRP 101 and FRRP 202.

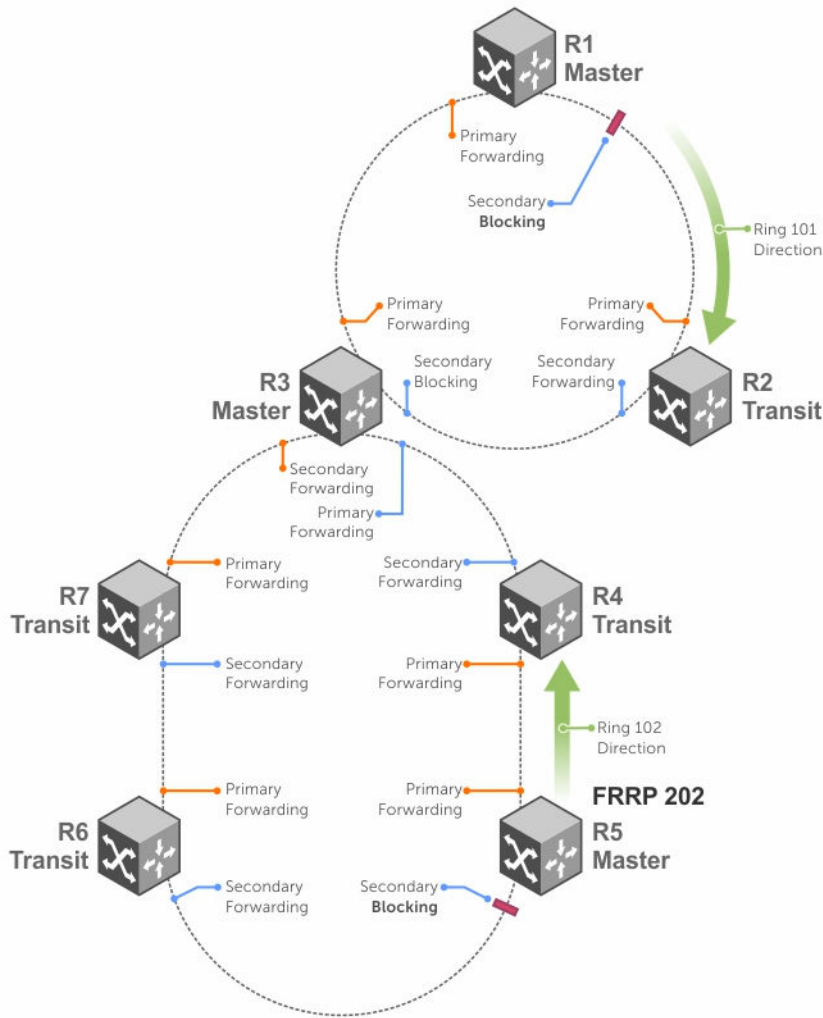


Figure 40. Multiple Rings Connected by a Single Switch Example

Important FRRP Points

FRRP provides a convergence time that can generally range between 150ms and 1500ms for Layer 2 networks.

The Master node originates a high-speed frame that circulates around the ring. This frame, appropriately, sets up or breaks down the ring.

- The Master node transmits ring status check frames at specified intervals.
- You can run multiple physical rings on the same switch.
- One Master node per ring — all other nodes are Transit.
- Each node has two member interfaces — primary and secondary.
- There is no limit to the number of nodes on a ring.
- Master node ring port states — blocking, pre-forwarding, forwarding, and disabled.

- Transit node ring port states — blocking, pre-forwarding, forwarding, and disabled.
- STP disabled on ring interfaces.
- Master node secondary port is in blocking state during Normal operation.
- Ring health frames (RHF)
 - Hello RHF: sent at 500ms (hello interval); Only the Master node transmits and processes these.
 - Topology Change RHF: triggered updates; processed at all nodes.

Important FRRP Concepts

The following table lists some important FRRP concepts.

Concept	Explanation
Ring ID	Each <i>ring</i> has a unique 8-bit ring ID through which the ring is identified (for example, FRRP 101 and FRRP 202).
Control VLAN	Each <i>ring</i> has a unique Control VLAN through which tagged ring health frames (RHF) are sent. Control VLANs are used only for sending RHF, and cannot be used for any other purpose.
Member VLAN	Each <i>ring</i> maintains a list of member VLANs. Member VLANs must be consistent across the entire ring.
Port Role	Each <i>node</i> has two ports for each ring: Primary and Secondary. The Master node Primary port generates RHF. The Master node Secondary port receives the RHF. On Transit nodes, there is no distinction between a Primary and Secondary interface when operating in the Normal state.
Ring Interface State	Each interface (port) that is part of the ring maintains one of four states” <ul style="list-style-type: none"> • Blocking State — Accepts ring protocol packets but blocks data packets. LLDP, FEFD, or other Layer 2 control packets are accepted. Only the Master node Secondary port can enter this state. • Pre-Forwarding State — A transition state before moving to the Forward state. Control traffic is forwarded but data traffic is blocked. The Master node Secondary port transitions through this state during ring bring-up. All ports transition through this state when a port comes up. • Pre-Forwarding State — A transition state before moving to the Forward state. Control traffic is forwarded but data traffic is blocked. The Master node Secondary port transitions through this state during ring bring-up. All ports transition through this state when a port comes up. • Disabled State — When the port is disabled or down, or is not on the VLAN.
Ring Protocol Timers	<ul style="list-style-type: none"> • Hello Interval — The interval when ring frames are generated from the Master node’s Primary interface (default 500 ms). The Hello interval is configurable in 50 ms increments from 50 ms to 2000 ms. • Dead Interval — The interval when data traffic is blocked on a port. The default is three times the Hello interval rate. The dead interval is configurable in 50 ms increments from 50 ms to 6000 ms.
Ring Status	The state of the FRRP ring. During initialization/configuration, the default ring status is Ring-down (disabled). The Primary and Secondary interfaces, control

Concept	Explanation
	<p>VLAN, and Master and Transit node information must be configured for the ring to be up.</p> <ul style="list-style-type: none"> • Ring-Up — Ring is up and operational. • Ring-Down — Ring is broken or not set up.
Ring Health-Check Frame (RHF)	<p>The Master node generates two types of RHFs. RHFs never loop the ring because they terminate at the Master node's secondary port.</p> <ul style="list-style-type: none"> • Hello RHF (HRHF) — These frames are processed only on the Master node's Secondary port. The Transit nodes pass the HRHF through without processing it. An HRHF is sent at every Hello interval. • Topology Change RHF (TCRHF) — These frames contains ring status, keepalive, and the control and member VLAN hash. The TCRHF is processed at each node of the ring. TCRHFs are sent out the Master Node's Primary and Secondary interface when the ring is declared in a Failed state with the same sequence number, on any topology change to ensure that all Transit nodes receive it. There is no periodic transmission of TCRHFs. The TCRHFs are sent on triggered events of ring failure or ring restoration only.

Implementing FRRP

- FRRP is media and speed independent.
- FRRP is a Dell proprietary protocol that does not interoperate with any other vendor.
- You must disable the spanning tree protocol (STP) on both the Primary and Secondary interfaces before you can enable FRRP.
- All ring ports must be Layer 2 ports. This is required for both Master and Transit nodes.
- A VLAN configured as a control VLAN for a ring cannot be configured as a control or member VLAN for any other ring.
- The control VLAN is not used to carry any data traffic; it carries only RHFs.
- The control VLAN cannot have members that are not ring ports.
- If multiple rings share one or more member VLANs, they cannot share any links between them.
- Member VLANs across multiple rings are not supported in Master nodes.
- Each ring has only one Master node; all others are transit nodes.

FRRP Configuration

These are the tasks to configure FRRP.

- [Creating the FRRP Group](#)
- [Configuring the Control VLAN](#)
 - Configure Primary and Secondary ports
- [Configuring and Adding the Member VLANs](#)

- Configure Primary and Secondary ports
- [Setting the FRRP Timers](#)

Other FRRP related commands are:

- [Clearing the FRRP Counters](#)
- [Viewing the FRRP Configuration](#)
- [Viewing the FRRP Information](#)

Creating the FRRP Group

Create the FRRP group on each switch in the ring. To create the FRRP group, use the command.

- Create the FRRP group with this Ring ID.
CONFIGURATION mode

```
protocol frrp ring-id
```

Ring ID: the range is from 1 to 255.

Configuring the Control VLAN

Control and member VLANs are configured normally for Layer 2. Their status as control or member is determined at the FRRP group commands.

For more information about configuring VLANs in Layer 2 mode, refer to [Layer 2](#).

Be sure to follow these guidelines:

- All VLANs must be in Layer 2 mode.
- You can only add ring nodes to the VLAN.
- A control VLAN can belong to one FRRP group only.
- Tag control VLAN ports.
- All ports on the ring must use the same VLAN ID for the control VLAN.
- You cannot configure a VLAN as both a control VLAN and member VLAN on the same ring.
- Only two interfaces can be members of a control VLAN (the Master Primary and Secondary ports).
- Member VLANs across multiple rings are not supported in Master nodes.

To create the control VLAN for this FRRP group, use the following commands on the switch that is to act as the Master node.

1. Create a VLAN with this ID number.
CONFIGURATION mode.

```
interface vlan vlan-id
```

VLAN ID: from 1 to 4094.

2. Tag the specified interface or range of interfaces to this VLAN.
CONFIG-INT-VLAN mode.

```
tagged interface slot/ port {range}
```

- For a 10/100/1000 Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information.
- For a Gigabit Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

Slot/Port, Range: Slot and Port ID for the interface. Range is entered Slot/Port-Port.

3. Assign the Primary and Secondary ports and the control VLAN for the ports on the ring.
CONFIG-FRRP mode.

```
interface primary int slot/port secondary int slot/port control-vlan vlan id
```

- For a 10/100/1000 Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information.
- For a Gigabit Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information.
- For a SONET interface, enter the keyword `sonet` then the slot/port information.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

Slot/Port, Range: Slot and Port ID for the interface. Range is entered Slot/Port-Port.

VLAN ID: The VLAN identification of the control VLAN.

4. Configure the Master node.
CONFIG-FRRP mode.

```
mode master
```

5. Identify the Member VLANs for this FRRP group.
CONFIG-FRRP mode.

```
member-vlan vlan-id {range}
```

VLAN-ID, Range: VLAN IDs for the ring's member VLANs.

6. Enable FRRP.
CONFIG-FRRP mode.

```
no disable
```

Configuring and Adding the Member VLANs

Control and member VLANs are configured normally for Layer 2. Their status as Control or Member is determined at the FRRP group commands.

For more information about configuring VLANs in Layer 2 mode, refer to the [Layer 2](#) chapter. Be sure to follow these guidelines:

- All VLANs must be in Layer 2 mode.

- Tag control VLAN ports. Member VLAN ports, except the Primary/Secondary interface, can be tagged or untagged.
- The control VLAN must be the same for all nodes on the ring.

To create the Members VLANs for this FRRP group, use the following commands on all of the Transit switches in the ring.

1. Create a VLAN with this ID number.
CONFIGURATION mode.

```
interface vlan vlan-id
```

VLAN ID: the range is from 1 to 4094.

2. Tag the specified interface or range of interfaces to this VLAN.
CONFIG-INT-VLAN mode.

```
tagged interface slot/port {range}
```

- For a 10/100/1000 Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information.
- For a Gigabit Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information.
- For a SONET interface, enter the keyword `sonet` then the slot/port information.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
- *Slot/Port, range*: Slot and Port ID for the interface. The range is entered Slot/Port-Port.

3. Assign the Primary and Secondary ports and the Control VLAN for the ports on the ring.
CONFIG-FRRP mode.

```
interface primary int slot/port secondary int slot/port control-vlan vlan id
```

- For a 10/100/1000 Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information.
- For a Gigabit Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information.
- For a SONET interface, enter the keyword `sonet` then the slot/port information.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

Slot/Port, Range: Slot and Port ID for the interface. Range is entered Slot/Port-Port.

VLAN ID: Identification number of the Control VLAN.

4. Configure a Transit node.
CONFIG-FRRP mode.

```
mode transit
```

5. Identify the Member VLANs for this FRRP group.

CONFIG-FRRP mode.

```
member-vlan vlan-id {range}
```

VLAN-ID, *Range*: VLAN IDs for the ring's Member VLANs.

6. Enable this FRRP group on this switch.

CONFIG-FRRP mode.

```
no disable
```

Setting the FRRP Timers

To set the FRRP timers, use the following command.

 **NOTE:** Set the Dead-Interval time 3 times the Hello-Interval.

- Enter the desired intervals for Hello-Interval or Dead-Interval times.
CONFIG-FRRP mode.

```
timer {hello-interval|dead-interval} milliseconds
```

- *Hello-Interval*: the range is from 50 to 2000, in increments of 50 (default is **500**).
- *Dead-Interval*: the range is from 50 to 6000, in increments of 50 (default is **1500**).

Clearing the FRRP Counters

To clear the FRRP counters, use one of the following commands.

- Clear the counters associated with this Ring ID.
EXEC PRIVELEGED mode.

```
clear frrp ring-id
```

Ring ID: the range is from 1 to 255.

- Clear the counters associated with all FRRP groups.
EXEC PRIVELEGED mode.

```
clear frrp
```

Viewing the FRRP Configuration

To view the configuration for the FRRP group, use the following command.

- Show the configuration for this FRRP group.
CONFIG-FRRP mode.

```
show configuration
```

Viewing the FRRP Information

To view general FRRP information, use one of the following commands.

- Show the information for the identified FRRP group.
EXEC or EXEC PRIVELEGED mode.

```
show frrp ring-id
```

Ring ID: the range is from 1 to 255.

- Show the state of all FRRP groups.
EXEC or EXEC PRIVELEGED mode.

```
show frrp summary
```

Ring ID: the range is from 1 to 255.

Troubleshooting FRRP

To troubleshoot FRRP, use the following information.

Configuration Checks

- Each Control Ring must use a unique VLAN ID.
- Only two interfaces on a switch can be Members of the same control VLAN.
- There can be only one Master node for any FRRP group.
- You can configure FRRP on Layer 2 interfaces only.
- Spanning Tree (if you enable it globally) must be disabled on both Primary and Secondary interfaces when you enable FRRP.
 - When the interface ceases to be a part of any FRRP process, if you enable Spanning Tree globally, also enable it explicitly for the interface.
- The maximum number of rings allowed on a chassis is 255.

Sample Configuration and Topology

The following example shows a basic FRRP topology.




Figure 41. Basic Topology and CLI Commands

Example of R1 MASTER

Example of R2 TRANSIT

Example of R3 TRANSIT

```
interface GigabitEthernet 1/24
 no ip address
 switchport
 no shutdown
!
interface GigabitEthernet 1/34
 no ip address
 switchport
 no shutdown
!
interface Vlan 101
 no ip address
 tagged GigabitEthernet 1/24,34
 no shutdown
!
interface Vlan 201
 no ip address
 tagged GigabitEthernet 1/24,34
 no shutdown
!
protocol frrp 101
```

```

interface primary GigabitEthernet 1/24
secondary GigabitEthernet 1/34 control-vlan 101
member-vlan 201
mode master
no disable

interface GigabitEthernet 2/14
no ip address
switchport
no shutdown
!
interface GigabitEthernet 2/31
no ip address
switchport
no shutdown
!
interface Vlan 101
no ip address
tagged GigabitEthernet 2/14,31
no shutdown
!
interface Vlan 201
no ip address
tagged GigabitEthernet 2/14,31
no shutdown
!
protocol frrp 101
interface primary GigabitEthernet 2/14 secondary GigabitEthernet 2/31 control-
vlan 101
member-vlan 201
mode transit
no disable

interface GigabitEthernet 3/14
no ip address
switchport
no shutdown
!
interface GigabitEthernet 3/21
no ip address
switchport
no shutdown
!
interface Vlan 101
no ip address
tagged GigabitEthernet 3/14,21
no shutdown
!
interface Vlan 201
no ip address
tagged GigabitEthernet 3/14,21
no shutdown

!
protocol frrp 101
interface primary GigabitEthernet 3/21
secondary GigabitEthernet 3/14 control-vlan 101
member-vlan 201
mode transit
no disable

```

GARP VLAN Registration Protocol (GVRP)

GARP VLAN registration protocol (GVRP) is supported on the MXL switch platform.

Typical virtual local area network (VLAN) implementation involves manually configuring each Layer 2 switch that participates in a given VLAN. GVRP, defined by the IEEE 802.1q specification, is a Layer 2 network protocol that provides for automatic VLAN configuration of switches. GVRP-compliant switches use GARP to register and de-register attribute values, such as VLAN IDs, with each other.

GVRP exchanges network VLAN information to allow switches to dynamically forward frames for one or more VLANs. Therefore, GVRP spreads this information and configures the needed VLANs on any additional switches in the network. Data propagates via the exchange of GVRP protocol data units (PDUs).

The purpose of GVRP is to simplify (but not eliminate) static configuration. The idea is to configure switches at the edge and have the information dynamically propagate into the core. As such, the edge ports must still be statically configured with VLAN membership information, and they do not run GVRP. It is this information that is propagated to create dynamic VLAN membership in the core of the network.

Important Points to Remember

- GVRP propagates VLAN membership throughout a network. GVRP allows end stations and switches to issue and revoke declarations relating to VLAN membership.
- VLAN registration is made in the context of the port that receives the GARP PDU and is propagated to the other active ports.
- GVRP is disabled by default; enable GVRP for the switch and then for individual ports.
- Dynamic VLANs are aged out after the LeaveAll timer expires three times without receipt of a Join message. To display status, use the `show gvrp statistics {interface interface | summary}` command.
- On the MXL Switch, you cannot enable per-VLAN spanning tree+ (PVST+) and GVRP at the same time. If spanning tree and GVRP are both required, implement either rapid spanning tree protocol (RSTP), spanning tree protocol (STP), or multiple spanning tree protocol (MSTP). The MXL 10/40GbE Switch IO Module system does support enabling GVRP and MSTP at the same time.

```
Dell(conf)#protocol spanning-tree pvst
Dell(conf-pvst)#no disable
% Error: GVRP running. Cannot enable PVST.
```

```
.....
Dell(conf)#protocol spanning-tree mstp
Dell(conf-mstp)#no disable
% Error: GVRP running. Cannot enable MSTP.
```

Configure GVRP

To begin, enable GVRP.

To facilitate GVRP communications, enable GVRP globally on each switch. Then, GVRP configuration is per interface on a switch-by-switch basis. Enable GVRP on each port that connects to a switch where you want GVRP information exchanged. In the following example, that type of port is referred to as a VLAN trunk port, but it is not necessary to specifically identify to the Dell Networking operating system (OS) that the port is a trunk port.




Figure 42. Global GVRP Configuration Example

Basic GVRP configuration is a two-step process:

1. [Enabling GVRP Globally](#)
2. [Enabling GVRP on a Layer 2 Interface](#)

Related Configuration Tasks

- [Configure GVRP Registration](#)
- [Configure a GARP Timer](#)

Enabling GVRP Globally

To configure GVRP globally, use the following command.

- Enable GVRP for the entire switch.
CONFIGURATION mode

```
gvrp enable
```

Example of Configuring GVRP

```
Dell(conf)#protocol gvrp
Dell(config-gvrp)#no disable
Dell(config-gvrp)#show config
!
protocol gvrp
  no disable
Dell(config-gvrp)#
```

To inspect the global configuration, use the `show gvrp brief` command.

Enabling GVRP on a Layer 2 Interface

To enable GVRP on a Layer 2 interface, use the following command.

- Enable GVRP on a Layer 2 interface.
INTERFACE mode

```
gvrp enable
```

Example of Enabling GVRP on an Interface

```
Dell(conf-if-gi-1/21)#switchport
Dell(conf-if-gi-1/21)#gvrp enable
Dell(conf-if-gi-1/21)#no shutdown
Dell(conf-if-gi-1/21)#show config
!
interface GigabitEthernet 1/21
  no ip address
  switchport
  gvrp enable
  no shutdown
```

To inspect the interface configuration, use the `show config` command from INTERFACE mode or use the `show gvrp interface` command in EXEC or EXEC Privilege mode.

Configure GVRP Registration

Configure GVRP registration.

There are three GVRP registration modes:

- **Normal Registration** — Allows dynamic creation, registration, and de-registration of VLANs (if you enabled dynamic VLAN creation). By default, the registration mode is set to Normal when you enable GVRP on a port. This default mode enables the port to dynamically register and de-register VLANs, and to propagate both dynamic and static VLAN information.
- **Fixed Registration Mode** — Configuring a port in fixed registration mode allows for manual creation and registration of VLANs, prevents VLAN deregistration, and registers all VLANs known on other ports on the port. For example, if an interface is statically configured via the CLI to belong to a VLAN, it should not be unconfigured when it receives a Leave PDU. Therefore, the registration mode on that interface is FIXED.
- **Forbidden Mode** — Disables the port to dynamically register VLANs and to propagate VLAN information except information about VLAN 1. A port with forbidden registration type thus allows only VLAN 1 to pass through even though the PDU carries information for more VLANs. Therefore, if you do not want the interface to advertise or learn about particular VLANs, set the interface to the registration mode of FORBIDDEN.

Based on the configuration in the following example, the interface 1/21 is not removed from VLAN 34 or VLAN 35 despite receiving a GVRP Leave message. Additionally, the interface is not dynamically added to VLAN 45 or VLAN 46, even if a GVRP Join message is received.

Example of the `gvrp registration` Command

```
Dell(conf-if-gi-1/21)#gvrp registration fixed 34,35
Dell(conf-if-gi-1/21)#gvrp registration forbidden 45,46
Dell(conf-if-gi-1/21)#show conf
!
interface GigabitEthernet 1/21
  no ip address
  switchport
  gvrp enable
  gvrp registration fixed 34-35
  gvrp registration forbidden 45-46
  no shutdown
Dell(conf-if-gi-1/21)#
```

Configure a GARP Timer

Set GARP timers to the same values on all devices that are exchanging information using GVRP.

There are three GARP timer settings.

- **Join** — A GARP device reliably transmits Join messages to other devices by sending each Join message two times. To define the interval between the two sending operations of each Join message, use this parameter. The default is **200ms**.
- **Leave** — When a GARP device expects to de-register a piece of attribute information, it sends out a Leave message and starts this timer. If a Join message does not arrive before the timer expires, the information is de-registered. The Leave timer must be greater than or equal to 3x the Join timer. The default is **600ms**.

- **LeaveAll** — After startup, a GARP device globally starts a LeaveAll timer. After expiration of this interval, it sends out a LeaveAll message so that other GARP devices can re-register all relevant attribute information. The device then restarts the LeaveAll timer to begin a new cycle. The LeaveAll timer must be greater than or equal to 5x of the Leave timer. The default is **10000ms**.

Example of the `garp timer` Command

```
Dell(conf)#garp timer leav 1000
Dell(conf)#garp timers leave-all 5000
Dell(conf)#garp timer join 300
```

Verification:

```
Dell(conf)#do show garp timer
GARP Timers Value (milliseconds)
-----
Join Timer      300
Leave Timer     1000
LeaveAll Timer   5000
Dell(conf)#
```

Dell displays this message if an attempt is made to configure an invalid GARP timer: `Dell(conf)#garp timers join 300 % Error: Leave timer should be >= 3*Join timer.`

Internet Group Management Protocol (IGMP)

Multicast is premised on identifying many hosts by a single destination IP address; hosts represented by the same IP address are a multicast group.

IGMP is a Layer 3 multicast protocol that hosts use to join or leave a multicast group. Multicast routing protocols (such as protocol-independent multicast [PIM]) use the information in IGMP messages to discover which groups are active and to populate the multicast routing table.

IGMP Protocol Overview

IGMP has three versions. Version 3 obsoletes and is backwards-compatible with version 2; version 2 obsoletes version 1.

IGMP Version 2

IGMP version 2 improves on version 1 by specifying IGMP Leave messages, which allows hosts to notify routers that they no longer care about traffic for a particular group.

Leave messages reduce the amount of time that the router takes to stop forwarding traffic for a group to a subnet (leave latency) after the last host leaves the group. In version 1 hosts quietly leave groups, and the router waits for a query response timer several times the value of the query interval to expire before it stops forwarding traffic.

To receive multicast traffic from a particular source, a host must join the multicast group to which the source is sending traffic. A host that is a member of a group is called a receiver. A host may join many groups, and may join or leave any group at any time. A host joins and leaves a multicast group by sending an IGMP message to its IGMP Querier. The querier is the router that surveys a subnet for multicast receivers and processes survey responses to populate the multicast routing table.

IGMP messages are encapsulated in IP packets, as shown in the following illustration.

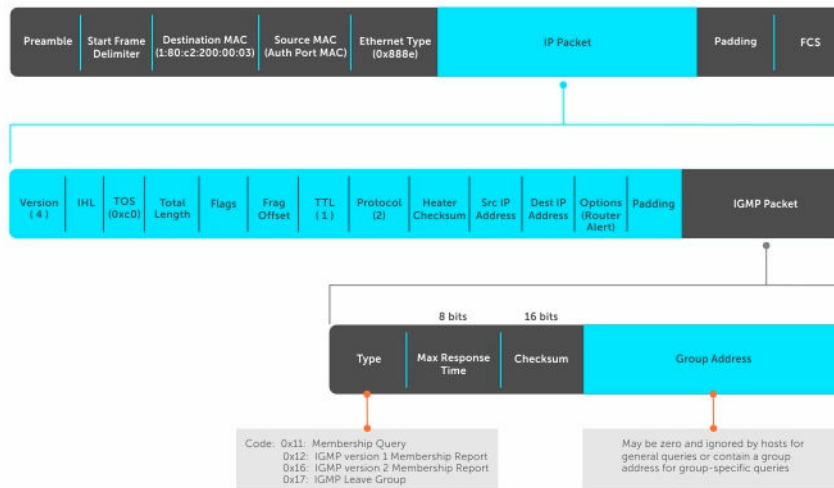


Figure 43. IGMP Messages in IP Packets

Join a Multicast Group

There are two ways that a host may join a multicast group: it may respond to a general query from its querier or it may send an unsolicited report to its querier.

- Responding to an IGMP Query
 - One router on a subnet is elected as the querier. The querier periodically multicasts (to all-multicast-systems address 224.0.0.1) a general query to all hosts on the subnet.
 - A host that wants to join a multicast group responds with an IGMP membership report that contains the multicast address of the group it wants to join (the packet is addressed to the same group). If multiple hosts want to join the same multicast group, only the report from the first host to respond reaches the querier, and the remaining hosts suppress their responses (for how the delay timer mechanism works, refer to [IGMP Snooping](#)).
 - The querier receives the report for a group and adds the group to the list of multicast groups associated with its outgoing port to the subnet. Multicast traffic for the group is then forwarded to that subnet.
- Sending an Unsolicited IGMP Report
 - A host does not have to wait for a general query to join a group. It may send an unsolicited IGMP membership report, also called an IGMP Join message, to the querier.

Leave a Multicast Group

The following describes how a host can leave a multicast group.

- A host sends a membership report of type 0x17 (IGMP Leave message) to the all routers multicast address 224.0.0.2 when it no longer cares about multicast traffic for a particular group.
- The querier sends a Group-Specific Query to determine whether there are any remaining hosts in the group. There must be at least one receiver in a group on a subnet for a router to forward multicast traffic for that group to the subnet.
- Any remaining hosts respond to the query according to the delay timer mechanism (refer to [IGMP Snooping](#)). If no hosts respond (because there are none remaining in the group), the querier waits a specified period and sends another query. If it still receives no response, the querier removes the

group from the list associated with forwarding port and stops forwarding traffic for that group to the subnet.

IGMP Version 3

Conceptually, IGMP version 3 behaves the same as version 2. However, there are differences.

- Version 3 adds the ability to filter by multicast source, which helps multicast routing protocols avoid forwarding traffic to subnets where there are no interested receivers.
- To enable filtering, routers must keep track of more state information, that is, the list of sources that must be filtered. An additional query type, the Group-and-Source-Specific Query, keeps track of state changes, while the Group-Specific and General queries still refresh the existing state.
- Reporting is more efficient and robust: hosts do not suppress query responses (non-suppression helps track state and enables the immediate-leave and IGMP snooping features), state-change reports are retransmitted to insure delivery, and a single membership report bundles multiple statements from a single host, rather than sending an individual packet for each statement.

The version 3 packet structure is different from version 2 to accommodate these protocol enhancements. Queries are still sent to the all-systems address 224.0.0.1, as shown in the following illustration, but reports are sent to the all IGMP version 3-capable multicast routers address 244.0.0.22, as shown in the second illustration.




Figure 44. IGMP Version 3 Packet Structure




Figure 45. IGMP Version 3–Capable Multicast Routers Address Structure

Joining and Filtering Groups and Sources

The following illustration shows how multicast routers maintain the group and source information from unsolicited reports.

1. The first unsolicited report from the host indicates that it wants to receive traffic for group 224.1.1.1.
2. The host's second report indicates that it is only interested in traffic from group 224.1.1.1, source 10.11.1.1. Include messages prevents traffic from all other sources in the group from reaching the subnet. Before recording this request, the querier sends a group-and-source query to verify that there are no hosts interested in any other sources. The multicast router must satisfy all hosts if they have conflicting requests. For example, if another host on the subnet is interested in traffic from 10.11.1.3, the router cannot record the include request. There are no other interested hosts, so the request is recorded. At this point, the multicast routing protocol prunes the tree to all but the specified sources.
3. The host's third message indicates that it is only interested in traffic from sources 10.11.1.1 and 10.11.1.2. Because this request again prevents all other sources from reaching the subnet, the router sends another group-and-source query so that it can satisfy all other hosts. There are no other interested hosts so the request is recorded.

Membership Reports: Joining and Filtering

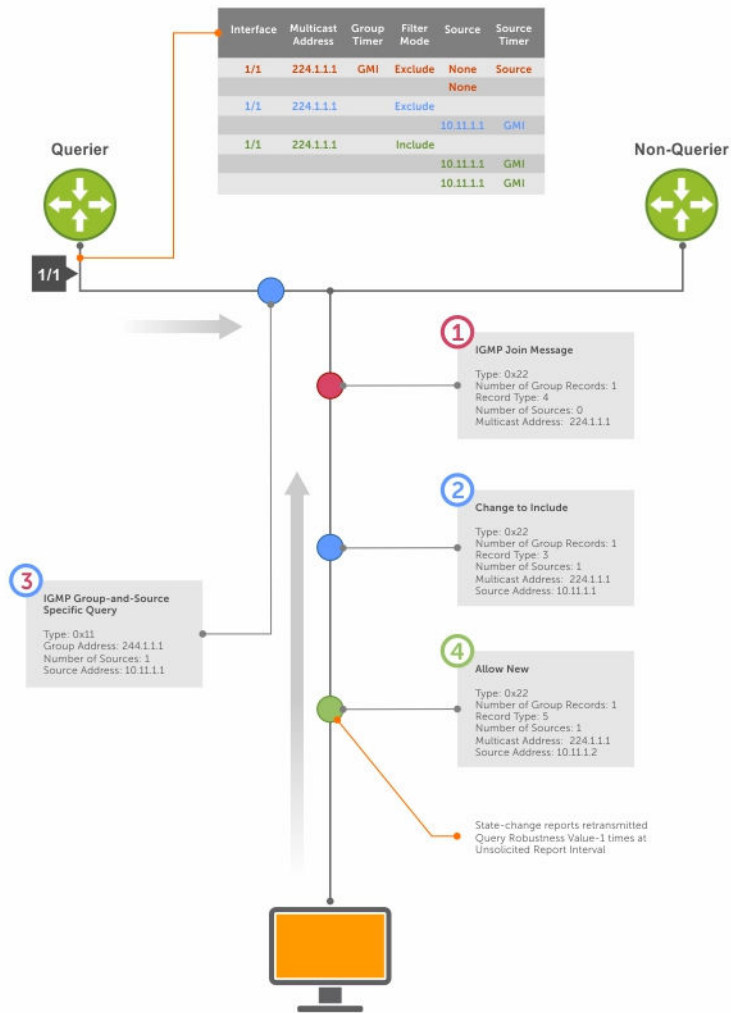


Figure 46. Membership Reports: Joining and Filtering

Leaving and Staying in Groups

The following illustration shows how multicast routers track and refresh state changes in response to group-and-specific and general queries.

- Host 1 sends a message indicating it is leaving group 224.1.1.1 and that the included filter for 10.11.1.1 and 10.11.1.2 are no longer necessary.
- The querier, before making any state changes, sends a group-and-source query to see if any other host is interested in these two sources; queries for state-changes are retransmitted multiple times. If any are, they respond with their current state information and the querier refreshes the relevant state information.
- Separately in the following illustration, the querier sends a general query to 224.0.0.1.
- Host 2 responds to the periodic general query so the querier refreshes the state information for that group.

Membership Reports: Joining and Filtering

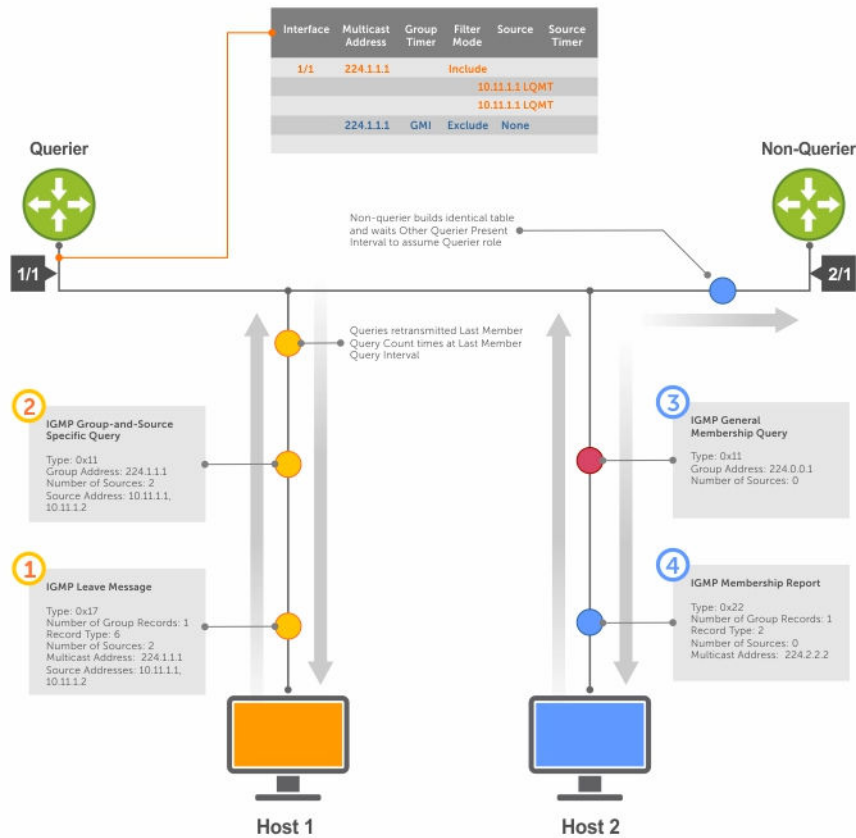


Figure 47. Membership Queries: Leaving and Staying

IGMP Snooping

IGMP snooping enables switches to use information in IGMP packets to generate a forwarding table that associates ports with multicast groups so that when they receive multicast frames, they can forward them only to interested receivers.

Multicast packets are addressed with multicast MAC addresses, which represent a group of devices, rather than one unique device. Switches forward multicast frames out of all ports in a virtual local area network (VLAN) by default, even though there may be only some interested hosts, which is a waste of bandwidth.

If you enable IGMP snooping on a VLT unit, IGMP snooping dynamically learned groups and multicast router ports are made to learn on the peer by explicitly tunneling the received IGMP control packets.

IGMP Snooping Implementation Information

- IGMP snooping on the Dell Networking OS uses IP multicast addresses not MAC addresses.
- IGMP snooping is not supported on stacked VLANs.

- IGMP snooping is supported on all MXL 10/40GbE stack members.
- IGMP snooping reacts to spanning tree protocol (STP) and multiple spanning tree protocol (MSTP) topology changes by sending a general query on the interface that transitions to the forwarding state.
-

Configuring IGMP Snooping

Configuring IGMP snooping is a one-step process. To enable, view, or disable IGMP snooping, use the following commands.

- Enable IGMP snooping on a switch.

CONFIGURATION mode

```
ip igmp snooping enable
```

- View the configuration.

CONFIGURATION mode

```
show running-config
```

- Disable snooping on a VLAN.

INTERFACE VLAN mode

```
no ip igmp snooping
```

Related Configuration Tasks

- [Enabling IGMP Immediate-Leave](#)
- [Disabling Multicast Flooding](#)
- [Specifying a Port as Connected to a Multicast Router](#)
- [Configuring the Switch as Querier](#)

Example of `ip igmp snooping enable` Command

```
Dell(conf)#ip igmp snooping enable
Dell(conf)#do show running-config igmp
ip igmp snooping enable
Dell(conf)#
```

Enabling IGMP Immediate-Leave

To remove a group-port association after receiving an IGMP Leave message, use the following command.

- Configure the switch to remove a group-port association after receiving an IGMP Leave message.

INTERFACE VLAN mode

```
ip igmp fast-leave
```

- View the configuration.

INTERFACE VLAN mode

```
show config
```

Example of the `show config` Command

```
Dell(conf-if-vl-100)#show config
!
interface Vlan 100
  no ip address
  ip igmp snooping fast-leave
  shutdown
Dell(conf-if-vl-100)#
```

Disabling Multicast Flooding

If the switch receives a multicast packet that has an IP address of a group it has not learned (unregistered frame), the switch floods that packet out of all ports on the VLAN.

On the MXL Switch, when you configure `no ip igmp snooping flood`, the system forwards the frames on the mrouter ports for first 96 IGMP snooping-enabled VLANs. For all other VLANs, the unregistered multicast packets are dropped.

Specifying a Port as Connected to a Multicast Router

To statically specify or view a port in a VLAN, use the following commands.

- Statically specify a port in a VLAN as connected to a multicast router.

```
INTERFACE VLAN mode
```

```
ip igmp snooping mrouter
```

- View the ports that are connected to multicast routers.

```
EXEC Privilege mode.
```

```
show ip igmp snooping mrouter
```

Configuring the Switch as Querier

To configure the switch as a querier, use the following command.

Hosts that do not support unsolicited reporting wait for a general query before sending a membership report. When the multicast source and receivers are in the same VLAN, multicast traffic is not routed and so there is no querier. Configure the switch to be the querier for a VLAN so that hosts send membership reports and the switch can generate a forwarding table by snooping.

- Configure the switch to be the querier for a VLAN by first assigning an IP address to the VLAN interface.

```
INTERFACE VLAN mode
```

```
ip igmp snooping querier
```

- IGMP snooping querier does not start if there is a statically configured multicast router interface in the VLAN.
- The switch may lose the querier election if it does not have the lowest IP address of all potential queriers on the subnet.
- When enabled, IGMP snooping querier starts after one query interval in case no IGMP general query (with IP SA lower than its VLAN IP address) is received on any of its VLAN members.

Adjusting the Last Member Query Interval

To adjust the last member query interval, use the following command.

When the querier receives a Leave message from a receiver, it sends a group-specific query out of the ports specified in the forwarding table. If no response is received, it sends another. The amount of time that the querier waits to receive a response to the initial query before sending a second one is the last member query interval (LMQI). The switch waits one LMQI after the second query before removing the group-port entry from the forwarding table.

- Adjust the last member query interval.

```
INTERFACE VLAN mode
```

```
ip igmp snooping last-member-query-interval
```

Fast Convergence after MSTP Topology Changes

The following describes the fast convergence feature.

When a port transitions to the Forwarding state as a result of an STP or MSTP topology change, the system sends a general query out of all ports except the multicast router ports. The host sends a response to the general query and the forwarding database is updated without having to wait for the query interval to expire.

When an IGMP snooping switch is not acting as a querier, it sends out the general query in response to the MSTP triggered link-layer topology change, with the source IP address of 0.0.0.0 to avoid triggering querier election.

Designating a Multicast Router Interface

To designate an interface as a multicast router interface, use the following command.

The Dell Networking OS also has the capability of listening in on the incoming IGMP general queries and designate those interfaces as the multicast router interface when the frames have a non-zero IP source address. All IGMP control packets and IP multicast data traffic originating from receivers is forwarded to multicast router interfaces.

- Designate an interface as a multicast router interface.

```
ip igmp snooping mrouter interface
```


Interfaces

This chapter describes 100/1000/10000 Mbps Ethernet, 10 Gigabit Ethernet, and 40 Gigabit Ethernet interface types, both physical and logical, and how to configure them with the Dell Networking operating software (OS).

Basic Interface Configuration

- [Interface Types](#)
- [View Basic Interface Information](#)
- [Enabling a Physical Interface](#)
- [Physical Interfaces](#)
- [Management Interfaces](#)
- [VLAN Interfaces](#)
- [Loopback Interfaces](#)
- [Null Interfaces](#)
- [Port Channel Interfaces](#)
- [Server Ports](#)

Advanced Interface Configuration

- [Bulk Configuration](#)
- [Define the Interface Range](#)
- [Monitoring and Maintaining Interfaces](#)
- [Splitting QSFP Ports to SFP+ Ports](#)
- [Configure MTU Size on an Interface](#)
- [Layer 2 Flow Control Using Ethernet Pause Frames](#)
- [Configure the MTU Size on an Interface](#)
- [Port-Pipes](#)
- [Auto-Negotiation on Ethernet Interfaces](#)
- [View Advanced Interface Information](#)

Interface Types

The following table describes different interface types.

Interface Type	Modes Possible	Default Mode	Requires Creation	Default State
Physical	L2, L3	Unset	No	Shutdown (disabled)
Management	N/A	N/A	No	No Shutdown (enabled)
Loopback	L3	L3	Yes	No Shutdown (enabled)
Null	N/A	N/A	No	Enabled
Port Channel	L2, L3	L3	Yes	Shutdown (disabled)
VLAN	L2, L3	L2	Yes (except default)	L2 - No Shutdown (enabled) L3 - Shutdown (disabled)

View Basic Interface Information

To view basic interface information, use the following command. You have several options for viewing interface status and configuration parameters.


- Lists all configurable interfaces on the chassis.


EXEC mode

```
show interfaces
```

This command has options to display the interface status, IP and MAC addresses, and multiple counters for the amount and type of traffic passing through the interface.

If you configured a port channel interface, this command lists the interfaces configured in the port channel.

 **NOTE:** To end output from the system, such as the output from the `show interfaces` command, enter `CTRL+C` and the system returns to the command prompt.

 **NOTE:** The CLI output may be incorrectly displayed as 0 (zero) for the Rx/Tx power values. To obtain the correct power information, perform a simple network management protocol (SNMP) query.

Example of the show interfaces Command

Example of the show ip interfaces brief Command

Example of the show running-config Command to View Physical Interfaces

The following example shows the configuration and status information for one interface.

```
Dell#show interfaces tengigabitethernet 0/16
TenGigabitEthernet 0/16 is up, line protocol is up
Hardware is DellForce10Eth, address is 00:1e:c9:f1:00:05
  Current address is 00:1e:c9:f1:00:05
Server Port AdminState is Up
Pluggable media not present
Interface index is 38080769
Internet address is not set
Mode of IP Address Assignment : NONE
DHCP Client-ID :tenG145001ec9f10005
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 10000 Mbit
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 5dlh18m
Queueing strategy: fifo
Input Statistics:
  34561 packets, 6266197 bytes
  38 64-byte pkts, 4373 over 64-byte pkts, 21491 over 127-byte pkts
  8659 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  21984 Multicasts, 12577 Broadcasts
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  44329 packets, 4722779 bytes, 0 underruns
  0 64-byte pkts, 44329 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  44329 Multicasts, 0 Broadcasts, 0 Unicasts
  0 throttles, 0 discarded, 0 collisions
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 4d0h28m
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  3 packets, 192 bytes, 0 underruns
  3 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  0 Multicasts, 3 Broadcasts, 0 Unicasts
  0 Vlans, 0 throttles, 0 discarded, 0 collisions
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:00:31
Dell
```

To view which interfaces are enabled for Layer 3 data transmission, use the show ip interfaces brief command in EXEC Privilege mode. In the following example, GigabitEthernet interface 1/5 is in Layer 3 mode because an IP address has been assigned to it and the interface's status is operationally up.

```
Dell#show ip interface brief
Interface          IP-Address  OK? Method  Status          Protocol
GigabitEthernet 1/0  unassigned NO  Manual  administratively down  down
GigabitEthernet 1/1  unassigned NO  Manual  administratively down  down
```

```
GigabitEthernet 1/2 unassigned YES Manual up up
GigabitEthernet 1/3 unassigned YES Manual up up
GigabitEthernet 1/4 unassigned YES Manual up up
GigabitEthernet 1/5 10.10.10.1 YES Manual up up
GigabitEthernet 1/6 unassigned NO Manual administratively down down
GigabitEthernet 1/7 unassigned NO Manual administratively down down
GigabitEthernet 1/8 unassigned NO Manual administratively down down
```

To view only configured interfaces, use the `show interfaces configured` command in the EXEC Privilege mode. In the previous example, GigabitEthernet interface 1/5 is in Layer 3 mode because an IP address has been assigned to it and the interface's status is operationally up.

To determine which physical interfaces are available, use the `show running-config` command in EXEC mode. This command displays all physical interfaces available on the line cards.

```
Dell#show running
Current Configuration ...
!
interface GigabitEthernet 9/6
  no ip address
  shutdown
!
interface GigabitEthernet 9/7
  no ip address
  shutdown
!
interface GigabitEthernet 9/8
  no ip address
  shutdown
!
interface GigabitEthernet 9/9
  no ip address
  shutdown
```

Enabling a Physical Interface

After determining the type of physical interfaces available, to enable and configure the interfaces, enter INTERFACE mode by using the `interface interface slot/port` command.

1. Enter the keyword `interface` then the type of interface and slot/port information.
CONFIGURATION mode

```
interface interface-type
```

- For the Management interface on the RPM, enter the keyword `ManagementEthernet` then the slot/port information.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

2. Enable the interface.
INTERFACE mode

```
no shutdown
```

To confirm that the interface is enabled, use the `show config` command in INTERFACE mode. To leave INTERFACE mode, use the `exit` command or `end` command. You cannot delete a physical interface.

Physical Interfaces

The switch interfaces support Layer 2 and Layer 3 traffic over the 100/1000/10000, 10-Gigabit, and 40-Gigabit Ethernet interfaces. These interfaces can also become part of virtual interfaces such as virtual local area networks (VLANs) or port channels.

For more information about VLANs, refer to [Bulk Configuration](#). For more information on port channels, refer to [Physical Interfaces](#).

Dell Networking OS Behavior: The MXL 10/40GbE switch systems use a single MAC address for all physical interfaces.

Configuration Task List for Physical Interfaces

By default, all interfaces are operationally disabled and traffic does not pass through them.

The following section includes information about optional configurations for physical interfaces:

- [Overview of Layer Modes](#)
- [Configuring Layer 2 \(Data Link\) Mode](#)
- [Configuring Layer 2 \(Interface\) Mode](#)
- [Configuring Layer 3 \(Interface\) Mode](#)
- [Configuring Layer 3 \(Network\) Mode](#)
- [Management Interfaces](#)
- [Auto-Negotiation on Ethernet Interfaces](#)
- [Adjusting the Keepalive Timer](#)
- [Clearing Interface Counters](#)

Overview of Layer Modes

On all systems running the Dell Networking OS, you can place physical interfaces, port channels, and VLANs in Layer 2 mode or Layer 3 mode.

By default, VLANs are in Layer 2 mode.

Type of Interface	Possible Modes	Requires Creation	Default State
10/100/1000 Ethernet, Gigabit Ethernet, 10 Gigabit Ethernet	Layer 2 Layer 3	No	Shutdown (disabled)
Management	N/A	No	Shutdown (disabled)
Loopback	Layer 3	Yes	No shutdown (enabled)
Null interface	N/A	No	Enabled

Type of Interface	Possible Modes	Requires Creation	Default State
Port Channel	Layer 2	Yes	Shutdown (disabled)
	Layer 3		
VLAN	Layer 2	Yes, except for the default VLAN.	No shutdown (active for Layer 2)
	Layer 3		Shutdown (disabled for Layer 3)

Configuring Layer 2 (Data Link) Mode

Do not configure switching or Layer 2 protocols such as spanning tree protocol (STP) on an interface unless the interface has been set to Layer 2 mode.

To set Layer 2 data transmissions through an individual interface, use the following command.

- Enable Layer 2 data transmissions through an individual interface.

```
INTERFACE mode
```

```
switchport
```

Example of a Basic Layer 2 Interface Configuration

```
Dell(conf-if)#show config
!
interface Port-channel 1
  no ip address
  switchport
  no shutdown
Dell(conf-if)#
```

Configuring Layer 2 (Interface) Mode

To configure an interface in Layer 2 mode, use the following commands.

- Enable the interface.

```
INTERFACE mode
```

```
no shutdown
```

- Place the interface in Layer 2 (switching) mode.

```
INTERFACE mode
```

```
switchport
```

For information about enabling and configuring the Spanning Tree Protocol, refer to [Spanning Tree Protocol \(STP\)](#).

To view the interfaces in Layer 2 mode, use the `show interfaces switchport` command in EXEC mode.

Configuring Layer 3 (Network) Mode

When you assign an IP address to a physical interface, you place it in Layer 3 mode.

To enable Layer 3 mode on an individual interface, use the following commands. In all interface types except VLANs, the `shutdown` command prevents all traffic from passing through the interface. In VLANs, the `shutdown` command prevents Layer 3 traffic from passing through the interface. Layer 2 traffic is unaffected by the `shutdown` command. One of the interfaces in the system must be in Layer 3 mode before you configure or enter a Layer 3 protocol mode (for example, OSPF).

- Enable Layer 3 on an individual interface

```
INTERFACE mode
```

```
ip address
```

- Enable the interface.

```
INTERFACE mode
```

```
no shutdown
```

```
Dell(conf-if)#show config
!
interface TenGigabitEthernet 1/5
  ip address 10.10.10.1 /24
  no shutdown
Dell(conf-if)#
```

Example of Error Due to Issuing a Layer 3 Command on a Layer 2 Interface

If an interface is in the incorrect layer mode for a given command, an error message displays (shown in bold). In the following example, the `ip address` command triggered an error message because the interface is in Layer 2 mode and the `ip address` command is a Layer 3 command only.

```
Dell(conf-if)#show config
!
interface GigabitEthernet 1/2
  no ip address
  switchport
  no shutdown
Dell(conf-if)#ip address 10.10.1.1 /24
% Error: Port is in Layer 2 mode Gi 1/2.
Dell(conf-if)#
```

To determine the configuration of an interface, use the `show config` command in INTERFACE mode or the various `show interface` commands in EXEC mode.

Configuring Layer 3 (Interface) Mode

To assign an IP address, use the following commands.

- Enable the interface.

```
INTERFACE mode
```

```
no shutdown
```

- Configure a primary IP address and mask on the interface.

INTERFACE mode

```
ip address ip-address mask [secondary]
```

The *ip-address* must be in dotted-decimal format (A.B.C.D) and the mask must be in slash format (/xx).

Add the keyword *secondary* if the IP address is the interface's backup IP address.

Example of the show ip interface Command

You can only configure one primary IP address per interface. You can configure up to 255 secondary IP addresses on a single interface.

To view all interfaces to see with an IP address assigned, use the `show ip interfaces brief` command in EXEC mode as shown in [View Basic Interface Information](#).


To view IP information on an interface in Layer 3 mode, use the `show ip interface` command in EXEC Privilege mode.

```
Dell(conf-if-vl-10)#do sh int vl 10
Vlan 10 is up, line protocol is up
Address is 00:1e:c9:f1:03:38, Current address is 00:1e:c9:f1:03:38
Interface index is 1107787786
Internet address is 5.5.5.1/24
Mode of IP Address Assignment : MANUAL
DHCP Client-ID: vlan10001ec9f10338
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 1000 Mbit
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:01:09
Queueing strategy: fifo
Time since last interface status change: 00:00:46
```

Management Interfaces

The IOM management interface has both a public IP and private IP address on the internal fabric D interface.

The public IP address is exposed to the outside world for Web GUI configurations/WSMAN and other proprietary traffic. You can statically configure the public IP address or obtain the IP address dynamically using the dynamic host configuration protocol (DHCP).

 **NOTE:** When you shut down a management interface, connectivity to the interface's private IP address is disabled.

You can access the full switch using:

- Internal RS-232 using the chassis management controller (CMC). Telnet into CMC and do a `connect -b switch-id` to get console access to corresponding IOM.
- External serial port with a universal serial bus (USB) connector (front panel): connect using the IOM front panel USB serial line to get console access (Labeled as USB B).
- Telnet/others using the public IP interface on the fabric D interface.
- CMC through the private IP interface on the fabric D interface.

The MXL switch system supports the management Ethernet interface as well as the standard interface on any front-end port. You can use either method to connect to the system.

Configuring Management Interfaces on the MXL Switch

On the MXL Switch IO Module, the dedicated management interface provides management access to the system.

You can configure this interface with the Dell Networking OS, but the configuration options on this interface are limited. You cannot configure Gateway addresses and IP addresses if it appears in the main routing table of the Dell Networking OS. In addition, proxy ARP is not supported on this interface.

To configure a management interface, use the following commands.

For additional management access, IOM supports the default VLAN (VLAN 1) L3 interface in addition to the public fabric D management interface. You can assign the IP address for the VLAN 1 default management interface using the setup wizard (or) through the CLI.

If you do not configure the VLAN 1 default using the wizard or CLI presented in startup-config, by default, the VLAN 1 management interface gets its IP address using DHCP.

There is only one management interface for the whole stack.

You can manage the MXL Switch from any port. Configure an IP address for the port using the `ip address` command. Enable the IP address for the port using the `no shutdown` command. You can use the `description` command from INTERFACE mode to note that the interface is the management interface. There is no separate management routing table, so you must configure all routes in the IP routing table (use the `ip route` command).

- Enter the slot and the port (0) to configure a Management interface.

CONFIGURATION mode

```
interface managementethernet interface
```

The slot range is 0–0.

- Configure an IP address and mask on a Management interface.

INTERFACE mode

```
ip address ip-address mask
```

– *ip-address mask*: enter an address in dotted-decimal format (A.B.C.D). The mask must be in / prefix format (/x).

Example of the `show interface` Command

To display the configuration for a given port, use the `show interface` command from EXEC Privilege mode, as shown in the following example.

To display the routing table for a given port, use the `show ip route` command from EXEC Privilege mode.

```
Dell#show int tengig 0/16
TenGigabitEthernet 0/16 is up, line protocol is down
Hardware is DellForce10Eth, address is 00:1e:c9:bb:02:c2
    Current address is 00:1e:c9:bb:02:c2
```

```


Server Port AdminState is Down
Pluggable media not present
Interface index is 38080769
Internet address is not set
Mode of IP Address Assignment : NONE
DHCP Client-ID :tenG145001ec9bb02c2
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 2w4d2h
Queueing strategy: fifo
Input Statistics:
  0 packets, 0 bytes
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  0 packets, 0 bytes, 0 underruns
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts, 0 Unicasts
  0 throttles, 0 discarded, 0 collisions
Rate info (interval 299 seconds):
Input 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Output 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 2w4d2h
Dell#


```

VLAN Interfaces

VLANs are logical interfaces and are, by default, in Layer 2 mode. Physical interfaces and port channels can be members of VLANs.


For more information about VLANs and Layer 2, refer to [Layer 2](#) and [Virtual LANs \(VLANs\)](#).

 **NOTE:** To monitor VLAN interfaces, use Management Information Base for Network Management of TCP/IP-based internets: MIB-II (RFC 1213).

 **NOTE:** You cannot simultaneously use egress rate shaping and ingress rate policing on the same VLAN.

The Dell Networking OS supports Inter-VLAN routing (Layer 3 routing in VLANs). You can add IP addresses to VLANs and use them in routing protocols in the same manner that physical interfaces are used. For more information about configuring different routing protocols, refer to the chapters on the specific protocol.

A consideration for including VLANs in routing protocols is that you must configure the `no shutdown` command. (For routing traffic to flow, you must enable the VLAN.)

 **NOTE:** You cannot assign an IP address to the default VLAN, which is VLAN 1 (by default). To assign another VLAN ID to the default VLAN, use the `default vlan-id vlan-id` command.

To assign an IP address to an interface, use the following command.

- Configure an IP address and mask on the interface.

INTERFACE mode

```
ip address ip-address mask [secondary]
```

- *ip-address mask*: enter an address in dotted-decimal format (A.B.C.D). The mask must be in slash format (/24).
- *secondary*: the IP address is the interface's backup IP address. You can configure up to eight secondary IP addresses.

Example of a Configuration for a VLAN Participating in an OSPF Process

```
interface Vlan 10
  ip address 1.1.1.2/24
  tagged GigabitEthernet 2/2-13
  tagged TenGigabitEthernet 5/0
  ip ospf authentication-key Dell force10
  ip ospf cost 1
  ip ospf dead-interval 60
  ip ospf hello-interval 15
  no shutdown
!
```

Loopback Interfaces

A Loopback interface is a virtual interface in which the software emulates an interface. Packets routed to it are processed locally.

Because this interface is not a physical interface, you can configure routing protocols on this interface to provide protocol stability. You can place Loopback interfaces in default Layer 3 mode.

To configure, view, or delete a Loopback interface, use the following commands.

- Enter a number as the Loopback interface.

CONFIGURATION mode

```
interface loopback number
```

The range is from 0 to 16383.

- View Loopback interface configurations.

EXEC mode

```
show interface loopback number
```

- Delete a Loopback interface.

CONFIGURATION mode

```
no interface loopback number
```

Many of the same commands found in the physical interface are also found in the Loopback interfaces.

For more information, refer to [Access Control Lists \(ACLs\)](#).

Null Interfaces

The Null interface is another virtual interface. There is only one Null interface. It is always up, but no traffic is transmitted through this interface.

To enter INTERFACE mode of the Null interface, use the following command.

- Enter INTERFACE mode of the Null interface.
CONFIGURATION mode

```
interface null 0
```

The only configurable command in INTERFACE mode of the Null interface is the `ip unreachable` command.

Port Channel Interfaces

Port channel interfaces support link aggregation, as described in IEEE Standard 802.3ad.

This section covers the following topics:

- [Port Channel Definition and Standards](#)
- [Port Channel Benefits](#)
- [Port Channel Implementation](#)
- [Configuration Tasks for Port Channel Interfaces](#)

Port Channel Definition and Standards

Link aggregation is defined by IEEE 802.3ad as a method of grouping multiple physical interfaces into a single logical interface—a link aggregation group (LAG) or port channel.

A LAG is “a group of links that appear to a MAC client as if they were a single link” according to IEEE 802.3ad. In the Dell Networking OS, a LAG is referred to as a port channel interface.

A port channel provides redundancy by aggregating physical interfaces into one logical interface. If one physical interface goes down in the port channel, another physical interface carries the traffic.

Port Channel Benefits

A port channel interface provides many benefits, including easy management, link redundancy, and sharing.

Port channels are transparent to network configurations and can be modified and managed as one interface. For example, you configure one IP address for the group and that IP address is used for all routed traffic on the port channel.

With this feature, you can create larger-capacity interfaces by utilizing a group of lower-speed links. For example, you can build a 40-Gigabit interface by aggregating four 10-Gigabit Ethernet interfaces together. If one of the five interfaces fails, traffic is redistributed across the three remaining interfaces.

Port Channel Implementation

The Dell Networking OS supports static and dynamic port channels.

- **Static** — Port channels that are statically configured.
- **Dynamic** — Port channels that are dynamically configured using the link aggregation control protocol (LACP). For details, refer to [Link Aggregation Control Protocol \(LACP\)](#).

There are 128 port-channels with 16 members per channel.

As soon as you configure a port channel, the system treats it like a physical interface. For example, IEEE 802.1Q tagging is maintained while the physical interface is in the port channel.

Member ports of a LAG are added and programmed into the hardware in a predictable order based on the port ID, instead of in the order in which the ports come up. With this implementation, load balancing yields predictable results across line card resets and chassis reloads.

A physical interface can belong to only one port channel at a time.

Each port channel must contain interfaces of the same interface type/speed.

Port channels can contain a mix of 100, 1000, or 10000 Mbps Ethernet interfaces and TenGigabit Ethernet interfaces. The interface speed (100, 1000, or 10000 Mbps) the port channel uses is determined by the first port channel member that is physically up. The system disables the interfaces that do not match the interface speed that the first channel member sets. That first interface may be the first interface that is physically brought up or was physically operating when interfaces were added to the port channel. For example, if the first operational interface in the port channel is a Gigabit Ethernet interface, all interfaces at 1000 Mbps are kept up, and all 100/1000/10000 interfaces that are not set to 1000 speed or auto negotiate are disabled.

100/1000/10000 Mbps Interfaces in Port Channels

When both 100/1000/10000 interfaces and TenGigabitEthernet interfaces are added to a port channel, the interfaces must share a common speed. When interfaces have a configured speed different from the port channel speed, the software disables those interfaces.

The common speed is determined when the port channel is first enabled. At that time, the software checks the first interface listed in the port channel configuration. If you enabled that interface, its speed configuration becomes the common speed of the port channel. If the other interfaces configured in that port channel are configured with a different speed, the system disables them.

For example, if four interfaces (TenGig 0/0, 0/1, 0/2, and 0/3) in which TenGig 0/0 and TenGig 0/3 are set to speed 100 Mb/s and the others are set to 10000 Mb/s, with all interfaces enabled, and you add them to a port channel by entering `channel-member tengigabitethernet 0/0-3` while in port channel interface mode, and the system determines if the first interface specified (TenGig 0/0) is up. After it is up, the common speed of the port channel is 100 Mb/s. The system disables those interfaces configured with speed 1000 Mb/s or whose speed is 1000 Mb/s as a result of auto-negotiation.

In this example, you can change the common speed of the port channel by changing its configuration so the first enabled interface referenced in the configuration is a 1000 Mb/s speed interface. You can also change the common speed of the port channel here by setting the speed of the TenGig 0/0 interface to 1000 Mb/s.

Configuration Tasks for Port Channel Interfaces

To configure a port channel (LAG), use the commands similar to those found in physical interfaces. By default, no port channels are configured in the startup configuration.

These are the mandatory and optional configuration tasks:

- [Creating a Port Channel](#) (mandatory)
- [Adding a Physical Interface to a Port Channel](#) (mandatory)
- [Reassigning an Interface to a New Port Channel](#) (optional)
- [Configuring the Minimum Oper Up Links in a Port Channel](#) (optional)
- [Adding or Removing a Port Channel from a VLAN](#) (optional)
- [Assigning an IP Address to a Port Channel](#) (optional)
- [Deleting or Disabling a Port Channel](#) (optional)

Creating a Port Channel

You can create up to 128 port channels with 16 port members per group on an MXL switch. To configure a port channel, use the following commands.

1. Create a port channel.
CONFIGURATION mode


```
interface port-channel id-number
```
2. Ensure that the port channel is active.
INTERFACE PORT-CHANNEL mode



```
no shutdown
```

After you enable the port channel, you can place it in Layer 2 or Layer 3 mode. To place the port channel in Layer 2 mode or configure an IP address to place the port channel in Layer 3 mode, use the `switchport` command.

You can configure a port channel as you would a physical interface by enabling or configuring protocols or assigning access control lists.


Adding a Physical Interface to a Port Channel

You can add any physical interface to a port channel if the interface configuration is minimal.

 **NOTE:** Port channels can contain a mix of 100/1000/10000 Ethernet interfaces and 10 Gigabit Ethernet interface, but the Dell Networking OS disables the interfaces that are not the same speed of the first channel member in the port channel (refer to [100/1000/10000 Mbps Interfaces in Port Channels](#)).

You can configure only the following commands on an interface if it is a member of a port channel:

- description
- shutdown/no shutdown
- mtu
- ip mtu (if the interface is on a Jumbo-enabled by default)

 **NOTE:** The MXL switch supports jumbo frames by default (the default maximum transmission unit [MTU] is 1554 bytes) You can configure the MTU using the `mtu` command from INTERFACE mode.

To view the interface's configuration, enter INTERFACE mode for that interface and use the `show config` command or from EXEC Privilege mode, use the `show running-config interface interface` command.

To add a physical interface to a port, use the following commands.

1. Add the interface to a port channel.
INTERFACE PORT-CHANNEL mode

```
channel-member interface
```

The `interface` variable is the physical interface type and slot/port information.

2. Double check that the interface was added to the port channel.
INTERFACE PORT-CHANNEL mode

```
show config
```

Example of the `show interfaces port-channel brief` Command

Example of the `show interface port-channel` Command

Example of Error Due to an Attempt to Configure an Interface that is Part of a Port Channel

To view the port channel's status and channel members in a tabular format, use the `show interfaces port-channel brief` command in EXEC Privilege mode, as shown in the following example.

```
Dell#show int port brief
Codes: L - LACP Port-channel

LAG Mode Status Uptime  Ports
1 L3 down 00:00:00 Te 0/16 (Down)
Dell#
```

The following example shows the port channel's mode (L2 for Layer 2 and L3 for Layer 3 and L2L3 for a Layer 2-port channel assigned to a routed VLAN), the status, and the number of interfaces belonging to the port channel.

```
Dell#show int port-channel
Port-channel 1 is down, line protocol is down
Hardware address is 00:1e:c9:f1:00:05, Current address is 00:1e:c9:f1:00:05
Interface index is 1107755009
Minimum number of links to bring Port-channel up is 1
Internet address is not set
Mode of IP Address Assignment : NONE
DHCP Client-ID :lag1001ec9f10005
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed auto
Members in this channel: Te 0/16(D)
ARP type: ARPA, ARP Timeout 04:00:00
```

```

Last clearing of "show interface" counters 00:05:44
Queueing strategy: fifo
Input Statistics:
  0 packets, 0 bytes
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  0 packets, 0 bytes, 0 underruns
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts, 0 Unicasts
  0 throttles, 0 discarded, 0 collisions
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:05:44

```

When more than one interface is added to a Layer 2-port channel, the system selects one of the active interfaces in the port channel to be the primary port. The primary port replies to flooding and sends protocol data units (PDUs). An asterisk in the `show interfaces port-channel brief` command indicates the primary port.

As soon as a physical interface is added to a port channel, the properties of the port channel determine the properties of the physical interface. The configuration and status of the port channel are also applied to the physical interfaces within the port channel. For example, if the port channel is in Layer 2 mode, you cannot add an IP address or a static MAC address to an interface that is part of that port channel.

In the following example, interface GigabitEthernet 1/6 is part of port channel 1, which is in Layer 2 mode, and an error message appeared when an IP address was configured.

```

Dell(conf-if-po-1)#show config
!
interface Port-channel 1
  no ip address
  channel-member TenGigabitEthernet 0/16
  shutdown
Dell(conf-if-po-1)#
Dell(conf-if-po-1)#int tengig 1/6
Dell(conf-if)#ip address 10.56.4.4 /24
% Error: Te 1/6 Port is part of a LAG.
Dell(conf-if)#

```

Reassigning an Interface to a New Port Channel

An interface can be a member of only one port channel. If the interface is a member of a port channel, remove it from the first port channel and then add it to the second port channel.

To reassign an interface to a new port channel, use the following commands.

1. Remove the interface from the first port channel.

```

INTERFACE PORT-CHANNEL mode

no channel-member interface

```
2. Change to the second port channel INTERFACE mode.

INTERFACE PORT-CHANNEL mode

```
interface port-channel id number
```

3. Add the interface to the second port channel.

INTERFACE PORT-CHANNEL mode

```
channel-member interface
```

Example of Moving an Interface to a New Port Channel

The following example shows moving the TenGigabitEthernet 1/8 interface from port channel 4 to port channel 3.

```
Dell(conf-if-po-1)#show config
!
interface Port-channel 1
  no ip address
  channel-member TenGigabitEthernet 0/16
  shutdown
Dell(conf-if-po-1)#no chann tengig 1/8
Dell(conf-if-po-1)#int port 5
Dell(conf-if-po-5)#channel tengig 1/8
Dell(conf-if-po-5)#show conf
!
interface Port-channel 5
  no ip address
  channel-member TenGigabitEthernet 1/8
  shutdown
Dell(conf-if-po-5)#
```

Configuring the Minimum Oper Up Links in a Port Channel

You can configure the minimum links in a port channel (LAG) that must be in “oper up” status to consider the port channel to be in “oper up” status.

To set the “oper up” status of your links, use the following command.

- Enter the number of links in a LAG that must be in “oper up” status.

INTERFACE mode

```
minimum-links number
```

The default is **1**.

Example of Configuring the Minimum Oper Up Links in a Port Channel

```
Dell#config t
Dell(conf)#int po 1
Dell(conf-if-po-1)#minimum-links 5
Dell(conf-if-po-1)#
```

Adding or Removing a Port Channel from a VLAN

As with other interfaces, you can add Layer 2 port channel interfaces to VLANs. To add a port channel to a VLAN, place the port channel in Layer 2 mode (by using the `switchport` command).

To add or remove a VLAN port channel and to view VLAN port channel members, use the following commands.

- Add the port channel to the VLAN as a tagged interface.

INTERFACE VLAN mode

```
tagged port-channel id number
```

An interface with tagging enabled can belong to multiple VLANs.

- Add the port channel to the VLAN as an untagged interface.

INTERFACE VLAN mode

```
untagged port-channel id number
```

An interface without tagging enabled can belong to only one VLAN.

- Remove the port channel with tagging enabled from the VLAN.

INTERFACE VLAN mode

```
no tagged port-channel id number
```

or

```
no untagged port-channel id number
```

- Identify which port channels are members of VLANs.

EXEC Privilege mode

```
show vlan
```

Assigning an IP Address to a Port Channel

You can assign an IP address to a port channel and use port channels in Layer 3 routing protocols. To assign an IP address, use the following command.

- Configure an IP address and mask on the interface.

INTERFACE mode

```
ip address ip-address mask [secondary]
```

- *ip-address mask*: enter an address in dotted-decimal format (A.B.C.D). The mask must be in slash format (/24).
- *secondary*: the IP address is the interface's backup IP address. You can configure up to eight secondary IP addresses.

Deleting or Disabling a Port Channel

To delete or disable a port channel, use the following commands.

- Delete a port channel.

CONFIGURATION mode

```
no interface portchannel channel-number
```

- Disable a port channel.

```
shutdown
```


When you disable a port channel, all interfaces within the port channel are operationally down also.

Load Balancing through Port Channels

Dell Networking OS uses hash algorithms for distributing traffic evenly over channel members in a port channel (LAG). The hash algorithm distributes traffic among ECMP paths and LAG members. The distribution is based on a flow, except for packet-based hashing. A flow is identified by the hash and is assigned to one link. In packet-based hashing, a single flow can be distributed on the LAG and uses one link.

Packet based hashing is used to load-balance traffic across a port-channel based on the IP Identifier field within the packet. Load balancing uses source and destination packet information to get the greatest advantage of resources by distributing traffic over multiple paths when transferring data to a destination.

Dell Networking OS allows you to modify the hashing algorithms used for flows and for fragments. The load-balance and hash-algorithm commands are available for modifying the distribution algorithms. Their syntax and implementation are somewhat different between the E-Series and the C-Series and S-Series

 **NOTE:** Hash-based load-balancing on MPLS does not work when packet-based hashing (load-balance ip-selection packet-based) is enabled.

Load-Balancing Method

For LAG hashing on the source IP, destination IP, source transmission control protocol (TCP)/user datagram protocol (UDP) port, and destination TCP/UDP port are used for hash computation by default. For packets without a Layer 3 header, Dell Networking OS automatically uses `load-balance mac source-dest-mac`.

Do not configure IP hashing or MAC hashing at the same time. If you configure an IP and MAC hashing scheme at the same time, the MAC hashing scheme takes precedence over the IP hashing scheme.

To change the IP traffic load-balancing default, use the following command.

- Replace the default IP 4-tuple method of balancing traffic over a port channel.

CONFIGURATION mode

```
[no] load-balance {ip-selection [dest-ip | source-ip]} | {mac [dest-mac | source-dest-mac | source-mac]} | {tcp-udp enable} {ipv6-selection} {tunnel} | {ingress-port}
```

You can select one, two, or all three of the following basic hash methods:

- `ip-selection [dest-ip | source-ip]` — Distribute IP traffic based on the IP destination or source address.
- `mac [dest-mac | source-dest-mac | source-mac]` — Distribute IPV4 traffic based on the destination or source MAC address, or both, along with the VLAN, Ethertype, source module ID and source port ID.
- `tcp-udp enable` — Distribute traffic based on the TCP/UDP source and destination ports.
- `ingress-port` — Option to Source Port Id for ECMP/ LAG hashing.

- `ipv6-selection`— Set the IPV6 key fields to use in hash computation.
- `tunnel`— Set the tunnel key fields to use in hash computation.

Hash Algorithm

The `load-balance` command selects the hash criteria applied to port channels.


If even distribution is not obtained with the `load-balance` command, the `hash-algorithm` command can be used to select the hash scheme for LAG, ECMP and NH-ECMP. The 12 bit Lag Hash can be rotated or shifted till the desired hash is achieved.

The `nh-ecmp` option allows you to change the hash value for recursive ECMP routes independently of non-recursive ECMP routes. This option provides for better traffic distribution over available equal cost links that involve a recursive next hop lookup.

Change the default (0) to another algorithm and apply it to ECMP, LAG hashing, or a particular line card.

CONFIGURATION mode

```
hash-algorithm {algorithm-number} | {ecmp {checksum|crc|xor} [number]} lag
{checksum|crc|xor} [number] nh-ecmp {[checksum|crc|xor} [number]}} | {linecard
number ip-sa-mask value ip-da-mask value}
```

 **NOTE:** To achieve the functionality of `hash-align` on the ExaScale platform, do not use CRC as an `hash-algorithm` method. For ExaScale systems, set the default `hash-algorithm` method to ensure CRC is not used for LAG. For example, `hash-algorithm ecmp xor lag checksum nh-ecmp checksum`

Example of hash algorithm

```
Dell(conf)##hash-algorithm ecmp xor 26 lag crc 26 nh-ecmp checksum 26
Dell(conf)#
```

The `hash-algorithm` command is specific to ECMP groups and has different defaults from the E-Series. The default ECMP hash configuration is `crc-lower`. This takes the lower 32 bits of the hash key to compute the egress port. Other options for ECMP hash-algorithms are:

- `crc-upper` — uses the upper 32 bits of the hash key to compute the egress port
- `dest-ip` — uses destination IP address as part of the hash key
- `lsb` — always uses the least significant bit of the hash key to compute the egress port

To change to another method, use the following command in the CONFIGURATION mode: `hash-algorithm ecmp {crc-upper} | {dest-ip} | {lsb}`

For more information on the algorithm choices, refer to the IP Routing chapter of the *Dell Networking OS Command Reference*.

Server Ports


By default, the MXL switch allows the server ports to come up as switch ports in `no shut` mode, ready to switch traffic.

Default Configuration without Start-up Config

This feature is enabled by default and can be enabled on reload by deleting the start-up config file.

On reload, all the server ports (1-32) come up as switch ports in No Shut mode. Uplinks remain in Shut mode ensuring that there are no network loops.

With this feature, you can install servers and test their connectivity by running applications on the servers, even before configuring VLAN membership, STP on all interfaces or uplinks.

 **NOTE:** This feature does not impact BMP mode. It always applies when reloading in Normal mode.

Important Points to Remember

- On a new MXL switch running the Dell Networking OS version 9.2(0.0), with no saved startup configuration, the switch comes up with all server ports as switch ports in No Shut state. When you configure STP, the switch brings up the uplink and saves the running configuration to the startup-config file. All the server ports without any specific configuration have the default configuration of Layer2 switch port and No Shut mode saved.
- On an existing MXL switch with a saved startup configuration, running an older Dell Networking OS version, an upgrade to a new version does not change the current behavior. This is because the start-up config file in older Dell Networking OS versions have the default configuration of Shut mode for all the server ports without any specific configuration. To enable this feature after upgrading a switch with a saved startup configuration, delete the start-up config file and reboot the switch. This allows all the server ports to come as Layer2 switch ports in No Shut state.
- In a stacked configuration of MXL switches, the behavior is similar to a standalone configuration. If a start-up config file is detected at bootup, the entire stack reboots using the saved configuration. If no start-up config file is detected at restart, the entire logical switch, including master unit, standby master, and any stack units restart with all server ports as Layer2 switch ports in No Shut mode.
- If a new stack unit is added to an existing stack, by default, the server side interfaces always start in Shut mode. If the startup configuration is deleted after a stack unit was added to a stack and the stack is reloaded, on reboot the entire logical switch comes up with all server ports as Layer2 switch ports in No Shut mode.

Bulk Configuration

Bulk configuration allows you to determine if interfaces are present for physical interfaces or configured for logical interfaces.


Interface Range


An interface range is a set of interfaces to which other commands may be applied and may be created if there is at least one valid interface within the range.

Bulk configuration excludes from configuration any non-existing interfaces from an interface range. A default VLAN may be configured only if the interface range being configured consists of only VLAN ports.

The `interface range` command allows you to create an interface range allowing other commands to be applied to that range of interfaces.

The interface range prompt offers the interface (with slot and port information) for valid interfaces. The maximum size of an interface range prompt is 32. If the prompt size exceeds this maximum, it displays (...) at the end of the output.

 **NOTE:** Non-existing interfaces are excluded from the interface range prompt.

 **NOTE:** When creating an interface range, interfaces appear in the order they were entered and are not sorted.

To display all interfaces that have been validated under the interface range context, use the `show range` command in Interface Range mode.

To display the running configuration only for interfaces that are part of interface range, use the `show configuration` command in Interface Range mode.

Bulk Configuration Examples

Use the `interface range` command for bulk configuration.

- [Create a Single-Range](#)
- [Create a Multiple-Range](#)
- [Exclude Duplicate Entries](#)
- [Exclude a Smaller Port Range](#)
- [Overlap Port Ranges](#)
- [Commas](#)
- [Add Ranges](#)

Create a Single-Range

The following is an example of a single range.

Example of the `interface range` Command (Single Range)

```
Dell(conf)# interface range tengigabitethernet 5/1 - 23
Dell(conf-if-range-te-5/1-23)# no shutdown
Dell(conf-if-range-te-5/1-23)#
```

Create a Multiple-Range

The following is an example of multiple range.

Example of the `interface range` Command (Multiple Ranges)

```
Dell(conf)#interface range tengigabitethernet 3/0 , tengigabitethernet 2/1 -
47 , vlan 1000
Dell(conf-if-range-te-2/1-47)#
```

Exclude Duplicate Entries

The following is an example showing how duplicate entries are omitted from the interface-range prompt.

Example of the Interface-Range Prompt for Duplicate Interfaces

```
Dell(conf)#interface range vlan 1 , vlan 1 , vlan 3 , vlan 3
Dell(conf-if-range-vl-1,vl-3)#
```

```
Dell(conf)#interface range tengigabitethernet 2/0 - 23 , tengigabitethernet 2/0
- 23 , tengigab 2/0 - 23
Dell(conf-if-range-te-2/0-23) #
```

Exclude a Smaller Port Range

The following is an example show how the smaller of two port ranges is omitted in the interface-range prompt.

Example of the Interface-Range Prompt for Multiple Port Ranges

```
Dell(conf)#interface range tengigabitethernet 2/0 - 23 , tengigab 2/1 - 10
Dell(conf-if-range-te-2/0-23) #
```

Overlap Port Ranges

The following is an example showing how the interface-range prompt extends a port range from the smallest start port number to the largest end port number when port ranges overlap. handles overlapping port ranges.

Example of the Interface-Range Prompt for Overlapping Port Ranges

```
Dell(conf)#inte ra tengig 2/1 - 11 , tengig 2/1 - 23
Dell(conf-if-range-te-2/1-23) #
```

Commas

The following is an example of how to use commas to add different interface types to the range, enabling all Gigabit Ethernet interfaces in the range 5/1 to 5/23 and both Ten Gigabit Ethernet interfaces 1/1 and 1/2.

Example of Multiple-Range Bulk Configuration Gigabit Ethernet and Ten-Gigabit Ethernet

```
Dell(conf-if)# interface range tengigabitethernet 5/1 - 23, tengigabitethernet
1/1 - 2
Dell(conf-if-range-te-5/1-23) # no shutdown
Dell(conf-if-range-te-5/1-23) #
```

Add Ranges

The following example shows how to use commas to add VLAN and port-channel interfaces to the range.

Example of Multiple-Range Bulk Configuration with VLAN and Port-channel

```
Dell(conf-ifrange-te-5/1-23-te-1/1-2) # interface range Vlan 2 - 100 , Port 1 -
25
Dell(conf-if-range-te-5/1-23-te-1/1-2-vl-2-100-po-1-25) # no shutdown
Dell(conf-if-range) #
```

Defining Interface Range Macros

You can define an interface-range macro to automatically select a range of interfaces for configuration. Before you can use the `macro` keyword in the `interface-range macro` command string, define the macro.

To define an interface-range macro, use the following command.

- Defines the interface-range macro and saves it in the running configuration file.

CONFIGURATION mode

```
define interface-range macro_name {vlan vlan_ID - vlan_ID} |
{{tengigabitethernet | fortyGigE} slot/ interface - interface} [ , {vlan
vlan_ID - vlan_ID} {{tengigabitethernet | fortyGigE} slot/interface -
interface}]
```

Define the Interface Range

The following example shows how to define an interface-range macro named “test” to select Fast Ethernet interfaces 5/1 through 5/4.

Example of the `define interface-range` Command for Macros

```
Dell(config)# define interface-range test tengigabitethernet 5/1 - 4
```

Choosing an Interface-Range Macro

To use an interface-range macro, use the following command.

- Selects the interfaces range to be configured using the values saved in a named interface-range macro.

CONFIGURATION mode

```
interface range macro name
```

Example of Using a Macro to Change the Interface Range Configuration Mode

The following example shows how to change to the interface-range configuration mode using the interface-range macro named “test.”

```
Dell(config)# interface range macro test
Dell(config-if)#
```

Monitoring and Maintaining Interfaces

Monitor interface statistics with the `monitor interface` command. This command displays an ongoing list of the interface status (up/down), number of packets, traffic statistics, and so on.

To view the interface’s statistics, use the following command.

- View the interface’s statistics.

EXEC Privilege mode

```
monitor interface interface
```

Enter the type of interface and slot/port information:

- For a 100/1000/10000 Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.

- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

Example of the `monitor interface` Command

The information displays in a continuous run, refreshing every 2 seconds by default. To manage the output, use the following keys.

- `m` – Change mode
- `l` – Page up
- `T` – Increase refresh interval (by 1 second)
- `t` – Decrease refresh interval (by 1 second)
- `c` – Clear screen
- `a` – Page down
- `q` – Quit

```
Dell#monitor interface tengig 3/1
```

```
Dell Networking uptime is 1 day(s), 4 hour(s), 31 minute(s)
Monitor time: 00:00:00 Refresh Intvl.: 2s
```

```
Interface: TenGig 3/1, Disabled, Link is Down, Linespeed is 1000 Mbit
```

```
Traffic statistics: Current      Rate      Delta
  Input bytes:         0      0 Bps      0
  Output bytes:        0      0 Bps      0
  Input packets:       0      0 pps      0
  Output packets:      0      0 pps      0
    64B packets:       0      0 pps      0
  Over 64B packets:    0      0 pps      0
  Over 127B packets:   0      0 pps      0
  Over 255B packets:   0      0 pps      0
  Over 511B packets:   0      0 pps      0
  Over 1023B packets:  0      0 pps      0
Error statistics:
  Input underruns:     0      0 pps      0
    Input giants:      0      0 pps      0
    Input throttles:   0      0 pps      0
    Input CRC:         0      0 pps      0
  Input IP checksum:   0      0 pps      0
    Input overrun:     0      0 pps      0
  Output underruns:    0      0 pps      0
  Output throttles:    0      0 pps      0
```

```
m - Change mode          c - Clear screen
l - Page up              a - Page down
T - Increase refresh interval  t - Decrease refresh interval
q - Quit
```

```
Dell
```


Maintenance Using TDR

The time domain reflectometer (TDR) is supported on all Dell Networking switch/routers.

TDR is an assistance tool to resolve link issues that helps detect obvious open or short conditions within any of the four copper pairs. TDR sends a signal onto the physical cable and examines the reflection of

the signal that returns. By examining the reflection, TDR is able to indicate whether there is a cable fault (when the cable is broken, becomes unterminated, or if a transceiver is unplugged).

TDR is useful for troubleshooting an interface that is not establishing a link; that is, when the link is flapping or not coming up. TDR is not intended to be used on an interface that is passing traffic. When a TDR test is run on a physical cable, it is important to shut down the port on the far end of the cable. Otherwise, it may lead to incorrect test results.

 **NOTE:** TDR is an intrusive test. Do not run TDR on a link that is up and passing traffic.

To test the condition of cables on 100/1000/10000 BASE-T modules, use the following commands.

1. To test for cable faults on the TenGigabitEthernet cable.

EXEC Privilege mode

```
tdr-cable-test tengigabitethernet <slot>/<port>
```

Between two ports, do not start the test on both ends of the cable.

Enable the interface before starting the test.

Enable the port to run the test or the test prints an error message.


2. Displays TDR test results.


EXEC Privilege mode

```
show tdr tengigabitethernet <slot>/<port>
```

Splitting QSFP Ports to SFP+ Ports

The MXL 10/40GbE switch supports splitting a 40GbE port on the base module or a 2-Port 40GbE QSFP+ module into four 10GbE SFP+ ports using a 4x10G breakout cable.

 **NOTE:** By default, the 40GbE ports on a 2-Port 40GbE QSFP+ module come up in 4x10GbE (quad) mode as eight 10GbE ports. On the base module, you must convert the 40GbE ports to 4x10GbE mode as described in the following section.

 **NOTE:** When you split a 40G port (such as fo 0/4) into four 10G ports, the 40G interface configuration is available in the startup configuration when you save the running configuration by using the write memory command. When a reload of the system occurs, the 40G interface configuration is not applicable because the 40G ports are split into four 10G ports after the reload operation. While the reload is in progress, you might see error messages when the configuration file is being loaded. You can ignore these error messages. Similarly, such error messages are displayed during a reload after you configure the four individual 10G ports to be stacked as a single 40G port.

- Split a single 40G port into 4-10G ports.

CONFIGURATION mode

```
stack-unit port number portmode quad
```

- `stack-unit`: Enter the stack member unit identifier of the stack member to reset. The range is from 0 to 5.
- `port <port number>`: Enter the port number of the 40G port to be split. The valid values on base module: 33 or 37; OPTM SLOT 0: 41 or 45; OPTM SLOT 1: 49 or 53.

- `portmode quad`: Identifies the uplink port as a split 10GbE SFP+ port.

To display the stack-unit number, enter the `show system brief` command.

- Save the configuration and reload the switch.

CONFIGURATION mode

```
write memory
```

```
reload
```

Merging SFP+ Ports to QSFP 40G Ports

To remove FANOUT mode in 40G QSFP Ports, use the following commands.

1. Merge 4-10G ports to a single 40G port.

CONFIGURATION mode

```
no stack-unit port number portmode quad
```

- *stack-unit*: Enter the stack member unit identifier of the stack member to reset. The range is from 0 to 5.
- `port <port number>`: Enter the port number of the 40GbE QSFP+ port. Valid values on base module: 33 or 37; OPTM SLOT 0: 41 or 45; OPTM SLOT 1: 49 or 53.
- `portmode quad`: Identifies the uplink port as a split 10GbE SFP+ port.

2. Save the configuration and reload the switch.

CONFIGURATION mode

```
write memory
```

```
reload
```

Important Points to Remember

- You cannot use split ports as stack-link to stack an MXL Switch.
- Split ports cannot be a part of any stacked system.
- The quad port must be in a default configuration before it can be split into 4x10G ports.
- The 40G port is lost in the configuration when the port is split; be sure the port is also removed from other L2/L3 feature configurations.
- The system must be reloaded after issuing the CLI for the change to take effect.

Configure the MTU Size on an Interface

The link MTU is the frame size of a packet. The IP MTU size is used for IP fragmentation.

If the system determines that the IP packet must be fragmented as it leaves the interface, the system divides the packet into fragments no bigger than the size set in the `ip mtu` command.

In the Dell Networking OS, MTU is defined as the entire Ethernet packet (Ethernet header + FCS + payload).

Because different networking vendors define MTU differently, check their documentation when planning MTU sizes across a network.

The following table lists the various Layer 2 overheads found in the Dell Networking OS and the number of bytes.

Table 24. Layer 2 Overhead

Transmission Media	MTU Range (in bytes)
Ethernet	594-12000 = link MTU 576-11982 = IP MTU

Converting a QSFP or QSFP+ Port to an SFP or SFP+ Port

You can convert a QSFP or QSFP+ port to an SFP or SFP+ port using the Quad to Small Form Factor Pluggable Adapter (QSA).


QSA provides smooth connectivity between devices that use Quad Lane Ports (such as the 40 Gigabit Ethernet adapters) and 10 Gigabit hardware that uses SFP+ based cabling. Using this adapter, you can effectively use a QSFP or QSFP+ module to connect to a lower-end switch or server that uses an SFP or SFP+ based module.


When connected to a QSFP or QSFP+ port on a 40 Gigabit adapter, QSA acts as an interface for the SFP or SFP+ cables. This interface enables you to directly plug in an SFP or SFP+ cable originating at a 10 Gigabit Ethernet port on a switch or server.

You can use QSFP optical cables (without a QSA) to split a 40 Gigabit port on a switch or a server into four 10 Gigabit ports. To split the ports, .enable the fan-out mode.

Similarly, you can enable the fan-out mode to configure the QSFP port on a device to act as an SFP or SFP+ port. As the QSA enables a QSFP or QSFP+ port to be used as an SFP or SFP+ port, Dell Networking OS does not immediately detect the QSA after you insert it into a QSFP port cage.

After you insert an SFP or SFP+ cable into a QSA connected to a 40 Gigabit port, Dell Networking OS assumes that all the four fanned-out 10 Gigabit ports have plugged-in SFP or SFP+ optical cables. However, the link UP event happens only for the first 10 Gigabit port and you can use only that port for data transfer. As a result, only the first fanned-out port is identified as the active 10 Gigabit port with a speed of 10G or 1G depending on whether you insert an SFP+ or SFP cable respectively.

 **NOTE:** Although it is possible to configure the remaining three 10 Gigabit ports, the Link UP event does not occur for these ports leaving the lanes unusable. Dell Networking OS perceives these ports to be in a Link Down state. You must not try to use these remaining three 10 Gigabit ports for actual data transfer or for any other related configurations.

 **NOTE:** Trident2 chip sets do not work at 1G speeds with auto-negotiation enabled. As a result, when you peer any device using SFP, the link does not come up if auto-negotiation is enabled. Therefore, disable auto-negotiation on platforms that currently use Trident2 chip sets (S6000 and Z9000). This limitation applies only when you convert QSFP to SFP using the QSA. This constraint does not apply for QSFP to SFP+ conversions using the QSA.

Important Points to Remember


- Before using the QSA to convert a 40 Gigabit Ethernet port to a 10 Gigabit SFP or SFP+ port, enable 40 G to 4*10 fan-out mode on the device.
- When you insert a QSA into a 40 Gigabit port, you can use only the first 10 Gigabit port in the fan-out mode to plug-in SFP or SFP+ cables. The remaining three 10 Gigabit ports are perceived to be in Link Down state and are unusable.
- You cannot use QSFP Optical cables on the same port where QSA is used.
- When you remove the QSA module alone from a 40 Gigabit port, without connecting any SFP or SFP+ cables; Dell Networking OS does not generate any event. However, when you remove a QSA module that has SFP or SFP+ optical cables plugged in, Dell Networking OS generates an SFP or SFP+ Removed event.
- The QSA module does not have a designated EEPROM. To recognize a QSA, Dell Networking OS reads the EEPROM corresponding to an SFP+ or SFP module that is plugged into QSA. The access location of this EEPROM is different from the EEPROM location of the QSFP+ module.
- The diagnostics application is capable of detecting insertion or removal of both the QSA as well as the SFP+ or SFP optical cables plugged into the QSA. In addition, the diagnostic application is also capable of reading the DDS and Vendor information from the EEPROM corresponding to SFP+ or SFP optical cables. As a result, no separate detection of QSA is required.

Example Scenarios

Consider the following scenarios:

- QSFP port 0 is connected to a QSA with SFP+ optical cables plugged in.
- QSFP port 4 is connected to a QSA with SFP optical cables plugged in.
- QSFP port 8 in fanned-out mode is plugged in with QSFP optical cables.
- QSFP port 12 in 40 G mode is plugged in with QSFP optical cables.

For these configurations, the following examples show the command output that the `show interfaces tengigbitethernet transceiver`, `show interfaces tengigbitethernet`, and `show inventory media` commands displays:

 **NOTE:** In the following `show interfaces tengigbitethernet` commands, the ports 1,2, and 3 are inactive and no physical SFP or SFP+ connection actually exists on these ports. However, Dell Networking OS still perceives these ports as valid and the output shows that pluggable media (optical cables) is inserted into these ports. This is a software limitation for this release.


Layer 2 Flow Control Using Ethernet Pause Frames

Ethernet pause frames allow for a temporary stop in data transmission.

A situation may arise where a sending device may transmit data faster than a destination device can accept it. The destination sends a pause frame back to the source, stopping the sender's transmission for a period of time.

The globally assigned 48-bit Multicast address 01-80-C2-00-00-01 is used to send and receive pause frames. To allow full duplex flow control, stations implementing the pause operation instruct the MAC to enable reception of frames with a destination address equal to this multicast address.

The pause frame is defined by IEEE 802.3x and uses MAC Control frames to carry the pause commands. Ethernet pause frames are supported on full duplex only. The only configuration applicable to half duplex ports is `rx off tx off`.

 **NOTE:** If a port is over-subscribed, Ethernet Pause Frame flow control does not ensure no loss behavior.

The following error message appears when trying to enable flow control when you already configured half duplex: `Can't configure flowcontrol when half duplex is configure, config ignored.`


The following error message appears when trying to enable half duplex and flow control configuration is on: `Can't configure half duplex when flowcontrol is on, config ignored.`

Enabling Pause Frames

Enable Ethernet pause frames flow control on all ports on a chassis. If not, the system may exhibit unpredictable behavior.

 **NOTE:** If you disable `rx flow control`, Dell Networking recommends rebooting the system.

The flow control sender and receiver must be on the same port-pipe. Flow control is not supported across different port-pipes. (also refer to [iSCSI Optimization: Operation](#)).

 **NOTE:** After you disable DCB, if link-level flow control is not automatically enabled on an interface, to enable flow control, manually shut down the interface (`shutdown` command) and re-enable it (`no shutdown` command).

To enable pause frames, use the following command.

- Control how the system responds to and generates 802.3x pause frames on 10 Gig ports.

INTERFACE mode

```
flowcontrol rx [off | on] tx [off | on] [negotiate]
```

- `rx on`: enter the keywords `rx on` to process the received flow control frames on this port.
- `rx off`: enter the keywords `rx off` to ignore the received flow control frames on this port.
- `tx on`: enter the keywords `tx on` to send control frames from this port to the connected device when a higher rate of traffic is received.
- `tx off`: enter the keywords `tx off` so that flow control frames are not sent from this port to the connected device when a higher rate of traffic is received.
- `negotiate`: enable pause-negotiation with the egress port of the peer device. If the `negotiate` command is not used, pause-negotiation is disabled. 40 gigabit Ethernet interfaces do not support pause-negotiation.

The default is **rx off**.

Configure MTU Size on an Interface

If a packet includes a Layer 2 header, the difference in bytes between the link MTU and IP MTU must be enough to include the Layer 2 header.

For example, for VLAN packets, if the IP MTU is 1400, the Link MTU must be no less than 1422:

1400-byte IP MTU + 22-byte VLAN Tag = 1422-byte link MTU

The MTU range is from 592 to 12000, with a default of 1500. IP MTU automatically configures.

The following table lists the various Layer 2 overheads found in the Dell Networking OS and the number of bytes.

Table 25. Layer 2 Overhead

Layer 2 Overhead	Difference Between Link MTU and IP MTU
Ethernet (untagged)	18 bytes
VLAN Tag	22 bytes
Untagged Packet with VLAN-Stack Header	22 bytes
Tagged Packet with VLAN-Stack Header	26 bytes

Link MTU and IP MTU considerations for port channels and VLANs are as follows.

Port Channels:

- All members must have the same link MTU value and the same IP MTU value.
- The port channel link MTU and IP MTU must be less than or equal to the link MTU and IP MTU values configured on the channel members.

For example, if the members have a link MTU of 2100 and an IP MTU 2000, the port channel's MTU values cannot be higher than 2100 for link MTU or 2000 bytes for IP MTU.

VLANs:

- All members of a VLAN must have the same IP MTU value.
- Members can have different Link MTU values. Tagged members must have a link MTU 4-bytes higher than untagged members to account for the packet tag.
- The VLAN link MTU and IP MTU must be less than or equal to the link MTU and IP MTU values configured on the VLAN members.

For example, the VLAN contains tagged members with Link MTU of 1522 and IP MTU of 1500 and untagged members with Link MTU of 1518 and IP MTU of 1500. The VLAN's Link MTU cannot be higher than 1518 bytes and its IP MTU cannot be higher than 1500 bytes.

Port-Pipes

A high-speed data bus connection used to switch traffic between front-end ports is known as the port pipe. A port pipe is a Dell Networking-specific term for the hardware path that packets follow through a system.


The MXL switch supports single port pipe only.

Auto-Negotiation on Ethernet Interfaces

By default, auto-negotiation of speed and duplex mode is enabled on 100/1000/10000 Base-T Ethernet interfaces. Only 10GE interfaces do not support auto-negotiation.

When using 10GE interfaces, verify that the settings on the connecting devices are set to no auto-negotiation.

The local interface and the directly connected remote interface must have the same setting, and auto-negotiation is the easiest way to accomplish that, as long as the remote interface is capable of auto-negotiation.

 **NOTE:** As a best practice, Dell Networking recommends keeping auto-negotiation enabled. Only disable auto-negotiation on switch ports that attach to devices not capable of supporting negotiation or where connectivity issues arise from interoperability issues.

For 100/1000/10000 Ethernet interfaces, the `negotiation auto` command is tied to the `speed` command. Auto-negotiation is always enabled when the `speed` command is set to 1000 in IOS.

Setting the Speed and Duplex Mode of Ethernet Interfaces

To discover whether the remote and local interface requires manual speed synchronization, and to manually synchronize them if necessary, use the following command sequence.

1. Determine the local interface status. Refer to the following example.

EXEC Privilege mode

```
show interfaces [interface] status
```

2. Determine the remote interface status.

EXEC mode or EXEC Privilege mode

[Use the command on the remote system that is equivalent to the first command.]

3. Access CONFIGURATION mode.

EXEC Privilege mode

```
config
```

4. Access the port.

CONFIGURATION mode

```
interface interface slot/port
```


5. Set the local port speed.

INTERFACE mode

```
speed {100 | 1000 | 10000 | auto}
```

6. Optionally, set full- or half-duplex.

INTERFACE mode

```
duplex {half | full}
```

7. Disable auto-negotiation on the port.

INTERFACE mode

```
no negotiation auto
```

If the speed was set to 1000, do not disable auto-negotiation.


8. Verify configuration changes.

INTERFACE mode

```
show config
```

Example of the `show interfaces status` Command to View Link Status

Example of Setting Port Speed and Disabling Auto-Negotiation

 **NOTE:** The `show interfaces status` command displays link status, but not administrative status. For both link and administrative status, use the `show ip interface [interface | brief] [configuration]` command.

```
Dell#show interfaces status
Port Description Status Speed Duplex Vlan
Te 0/1          Down   Auto   Auto   --
Te 0/2          Down   Auto   Auto   --
Te 0/3          Down   Auto   Auto   --
Te 0/4          Down   Auto   Auto   --
Te 0/5          Down   Auto   Auto   --
Te 0/6          Down   Auto   Auto   --
Te 0/7          Down   Auto   Auto   --
Te 0/8          Down   Auto   Auto   --
Te 0/9          Down   Auto   Auto   --
Te 0/10         Down   Auto   Auto   --
Te 0/11         Down   Auto   Auto   --
Te 0/12         Down   Auto   Auto   --
Te 0/13         Down   Auto   Auto   --
[output omitted]
```

In the previous example, several ports display "Auto" in the Speed field, including port 0/1. In the following example, the speed of port 0/1 is set to 100Mb and then its auto-negotiation is disabled.

```
Dell#configure
Dell(config)#interface tengig 0/1
Dell(conf-if-te-0/1)#speed 100
Dell(conf-if-te-0/1)#duplex full
Dell(conf-if-te-0/1)#no negotiation auto
Dell(conf-if-te-0/1)#show config
!
interface TenGigabitEthernet 0/1
no ip address
```

```
speed 100
duplex full
no shutdown
```

Set Auto-Negotiation Options

The `negotiation auto` command provides a mode option for configuring an individual port to forced master/ forced slave after you enable auto-negotiation.

⚠ CAUTION: Ensure that only one end of the node is configured as forced-master and the other is configured as forced-slave. If both are configured the same (that is, both as forced-master or both as forced-slave), the `show interface` command flaps between an auto-neg-error and forced-master/slave states.

Example of the `negotiation auto` Command

```
Dell(conf)# int tengig 0/0
Dell(conf-if)#neg auto
Dell(conf-if-autoneg)# ?

end                Exit from configuration mode
exit              Exit from autoneg configuration mode
mode             Specify autoneg mode
no               Negate a command or set its defaults
show            Show autoneg configuration information
Dell(conf-if-autoneg)#mode ?
forced-master    Force port to master mode
forced-slave     Force port to slave mode
Dell(conf-if-autoneg)#
```

Adjusting the Keepalive Timer

To change the time interval between keepalive messages on the interfaces, use the `keepalive` command. The interface sends keepalive messages to itself to test network connectivity on the interface. To change the default time interval between keepalive messages, use the following command.

- Change the default interval between keepalive messages.

```
INTERFACE mode
```

```
keepalive [seconds]
```

- View the new setting.

```
INTERFACE mode
```

```
show config
```

View Advanced Interface Information

The following options have been implemented for the `show [ip | running-config] interfaces` commands.

When you use the `configured` keyword, only interfaces that have non-default configurations display.

Example of `show` Commands that Use the `configured` Keyword

Example of the `show interfaces switchport` Command

The following example lists the possible `show` commands that have the `configured` keyword available:

```

Dell#show interfaces configured
Dell#show interfaces tengigabitEthernet 0 configured
Dell#show ip interface configured
Dell#show ip interface tengigabitEthernet 1 configured
Dell#show interfaces fortygigabitEthernet 0 configured
Dell#show ip interface fortygigabitEthernet 1 configured
Dell#show ip interface brief configured
Dell#show running-config interfaces configured
Dell#show running-config interface tengigabitEthernet 1 configured

```

In EXEC mode, the `show interfaces switchport` command displays only interfaces in Layer 2 mode and their relevant configuration information. The `show interfaces switchport` command displays the interface, whether it supports IEEE 802.1Q tagging or not, and the VLANs to which the interface belongs.

```

Dell#show interfaces switchport
Name: TenGigabitEthernet 13/0
802.1QTagged: True
Vlan membership:
Vlan 2

Name: TenGigabitEthernet 13/1
802.1QTagged: True
Vlan membership:
Vlan 2

Name: TenGigabitEthernet 13/2
802.1QTagged: True
Vlan membership:
Vlan 2

Name: TenGigabitEthernet 13/3
802.1QTagged: True
Vlan membership:
Vlan 2
--More--

```

Configuring the Interface Sampling Size

You can enter any value between five and 299 seconds (the default). If you enter 1 to 5 seconds, software polling is done at 5 second intervals. If you enter 6 to 10 seconds, software polling is done at 10 second interval. For any other value, software polling is done once every 15 seconds. So, for example, if you enter "19", you actually get a sample of the past 15 seconds.

All LAG members inherit the rate interval configuration from the LAG.

The following example shows how to configure rate interval when changing the default value.

To configure the number of seconds of traffic statistics to display in the `show interfaces` output, use the following command.

- Configure the number of seconds of traffic statistics to display in the `show interfaces` output.

INTERFACE mode

```
rate-interval
```

Example of the `rate-interval` Command

The bold lines shows the default value of 299 seconds, the change-rate interval of 100, and the new rate interval set to 100.

```
Dell#show interfaces
TenGigabitEthernet 10/0 is down, line protocol is down
Hardware is Dell Force10Eth, address is 00:01:e8:01:9e:d9
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 10000 Mbit
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 1d23h44m
Queueing strategy: fifo
  0 packets input, 0 bytes
  Input 0 IP Packets, 0 Vlans 0 MPLS
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  Received 0 input symbol errors, 0 runts, 0 giants, 0 throttles
  0 CRC, 0 IP Checksum, 0 overrun, 0 discarded
  0 packets output, 0 bytes, 0 underruns
  Output 0 Multicasts, 0 Broadcasts, 0 Unicasts
  0 IP Packets, 0 Vlans, 0 MPLS
  0 throttles, 0 discarded
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 1d23h40m
```


```
Dell(conf)#interface tengigabitethernet 10/0
Dell(conf-if-te-10/0)#rate-interval 100
```

```
Dell#show interfaces
TenGigabitEthernet 10/0 is down, line protocol is down
Hardware is Dell Force10Eth, address is 00:01:e8:01:9e:d9
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 10000 Mbit
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 1d23h45m
Queueing strategy: fifo
  0 packets input, 0 bytes
  Input 0 IP Packets, 0 Vlans 0 MPLS
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  Received 0 input symbol errors, 0 runts, 0 giants, 0 throttles
  0 CRC, 0 IP Checksum, 0 overrun, 0 discarded
  0 packets output, 0 bytes, 0 underruns
  Output 0 Multicasts, 0 Broadcasts, 0 Unicasts
  0 throttles, 0 discarded
Rate info (interval 100 seconds):
  Input 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Time since last interface status change: 1d23h42m
```

Dynamic Counters

By default, counting is enabled for IPFLOW, IPACL, L2ACL, L2FIB.

For the remaining applications, the system automatically turns on counting when you enable the application, and is turned off when you disable the application.

 **NOTE:** If you enable more than four counter-dependent applications on a port pipe, there is an impact on line rate performance.

The following counter-dependent applications are supported by the Dell Networking OS:

- Egress VLAN
- Ingress VLAN
- Next Hop 2
- Next Hop 1
- Egress ACLs
- ILM
- IP FLOW
- IP ACL
- IP FIB
- L2 ACL
- L2 FIB

Clearing Interface Counters

The counters in the `show interfaces` command are reset by the `clear counters` command. This command does not clear the counters any SNMP program captures.

To clear the counters, use the following the command.

- Clear the counters used in the `show interface` commands for all VRRP groups, VLANs, and physical interfaces or selected ones. Without an interface specified, the command clears all interface counters.
EXEC Privilege mode

```
clear counters [interface] [vrrp [vrid] | learning-limit]
```

(OPTIONAL) Enter the following interface keywords and slot/port or number information:

- For a Loopback interface, enter the keyword `loopback` then a number from 0 to 16383.
- For a Port Channel interface, enter the keywords `port-channel` then a number from 1 to 128.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
- For a VLAN, enter the keyword `vlan` then a number from 1 to 4094.
- (OPTIONAL) To clear statistics for all VRRP groups configured, enter the keyword `vrrp`. Enter a number from 1 to 255 as the `vrid`.
- (OPTIONAL) To clear unknown source address (SA) drop counters when you configure the MAC learning limit on the interface, enter the keywords `learning-limit`.

Example of the `clear counters` Command

When you enter this command, confirm that you want the Dell Networking OS to clear the interface counters for that interface.

```
Dell#clear counters tengig 0/0  
Clear counters on TenGigabitEthernet 0/0 [confirm]  
Dell#
```

Enhanced Validation of Interface Ranges

You can avoid specifying spaces between the range of interfaces, separated by commas, that you configure by using the `interface range` command. For example, if you enter a list of interface ranges, such as `interface range fo 2/0-1,te 10/0,gi 3/0,fa 0/0`, this configuration is considered valid. The comma-separated list is not required to be separated by spaces in between the ranges. You can associate multicast MAC or hardware addresses to an interface range and VLANs by using the `mac-address-table static multicast-mac-address vlan vlan-id output-range interface` command.

Enhanced Control of Remote Fault Indication Processing

By default, MXL module processes RFI errors transmitted by remote peers and brings down the interface when an RFI error is detected.

You must enter the interface configuration mode before configuring Remote Fault Indication (RFI).

You can use the following CLI commands to enable or disable processing of received RFI events:

```
Dell(conf-if-te-1/3)#remote-fault-signaling rx ?  
on Enable  
off Disable
```


The default is "remote-fault-signaling rx on".

Internet Protocol Security (IPSec)

IPSec is an end-to-end security scheme for protecting IP communications by authenticating and encrypting all packets in a communication session.

Use IPSec between hosts, between gateways, or between hosts and gateways.

IPSec is compatible with Telnet and file transfer protocols (FTPs) and can operate in Transport mode. In Transport mode, IPSec encrypts only the packet payload; the IP header is unchanged. This is the default mode.

 **NOTE:** Due to performance limitations on the control processor, you cannot enable IPSec on all packets in a communication session.

IPSec uses the following protocols:

- **Authentication Headers (AH)** — Disconnected integrity and origin authentication for IP packets
- **Encapsulating Security (ESP)** — Confidentiality, authentication, and data integrity for IP packets
- **Security Associations (SA)** — Necessary algorithmic parameters for AH and ESP functionality

IPSec supports the following authentication and encryption algorithms:

- Authentication only:
 - MD5
 - SHA1
- Encryption only:
 - 3DES
 - CBC
 - DES
- ESP Authentication and Encryption:
 - MD5 and 3DES
 - MD5 and CBC
 - MD5 and DES
 - SHA1 and 3DES
 - SHA1 and CBC
 - SHA1 and DES

Configuring IPSec

The following sample configuration shows how to configure FTP and telnet for IPSec.

1. Define the transform set.

CONFIGURATION mode

```
crypto ipsec transform-set myXform-seta esp-authentication md5 esp-  
encryption des
```

2. Define the crypto policy.

CONFIGURATION mode

```
crypto ipsec policy
```

```
myCryptoPolicy 10 ipsec-manual
```

```
transform-set myXform-set
```

```
session-key inbound esp 256
```

```
auth <key> encrypt <key>
```

```
session-key outbound esp 257
```

```
auth <key> encrypt <key>
```

```
match 0 tcp a::1 /128 0 a::2 /128 21
```

```
match 1 tcp a::1 /128 21 a::2 /128 0
```

```
match 2 tcp 1.1.1.1 /32 0 1.1.1.2 /32 21
```

```
match 3 tcp 1.1.1.1 /32 21 1.1.1.2 /32 0
```

3. Apply the crypto policy to management traffic.

CONFIGURATION mode

```
management crypto-policy
```

```
myCryptoPolicy
```


IPv4 Routing

The Dell Networking OS supports various IP addressing features.

This chapter describes the basics of domain name service (DNS), address resolution protocol (ARP), and routing principles and their implementation in the Dell Networking operating system (OS).

IP Feature	Default
DNS	Disabled
Directed Broadcast	Disabled
Proxy ARP	Enabled
ICMP Unreachable	Disabled
ICMP Redirect	Disabled

IP Addresses

The Dell Networking OS supports IP version 4, as described in RFC 791. It also supports classful routing and variable length subnet masks (VLSM).


With VLSM, you can configure one network with different masks. Supernetting, which increases the number of subnets, is also supported. To subnet, you add a mask to the IP address to separate the network and host portions of the IP address.


At its most basic level, an IP address is 32-bits composed of network and host portions and represented in dotted decimal format. For example, 00001010110101100101011110000011 is represented as 10.214.87.131.

For more information about IP addressing, refer to RFC 791, *Internet Protocol*.

Implementation Information

In the Dell Networking OS, you can configure any IP address as a static route except IP addresses already assigned to interfaces.

 **NOTE:** The Dell Networking OS versions 7.7.1.0 and later support 31-bit subnet masks (/31, or 255.255.255.254) as defined by RFC 3021. This feature allows you to save two more IP addresses on point-to-point links than 30-bit masks. The system supports RFC 3021 with ARP.

 **NOTE:** Even though Dell Networking OS listens to all ports, you can only use the ports starting from 35001 for IPv4 traffic. Ports starting from 0 to 35000 are reserved for internal use and you cannot use them for IPv4 traffic.

Configuration Tasks for IP Addresses

The following describes the tasks associated with IP address configuration.

Configuration tasks for IP addresses includes:

- [Assigning IP Addresses to an Interface](#) (mandatory)
- [Configuring Static Routes](#) (optional)
- [Configure Static Routes for the Management Interface](#) (optional)

For a complete listing of all commands related to IP addressing, refer to the *Dell Networking OS Command Line Interface Reference Guide*.

Assigning IP Addresses to an Interface

Assign primary and secondary IP addresses to physical or logical (for example, [virtual local area network [VLAN] or port channel) interfaces to enable IP communication between the system and hosts connected to that interface.

In the system, you can assign one primary address and up to 255 secondary IP addresses to each interface.

1. Enter the keyword `interface` then the type of interface and slot/port information.

CONFIGURATION mode

```
interface interface
```

- For a Loopback interface, enter the keyword `loopback` then a number from 0 to 16383.
- For the Management interface on the RPM, enter the keyword `ManagementEthernet` then the slot/port information. The slot range is from 0 to 1. The port range is 0/0.
- For a port channel interface, enter the keywords `port-channel` then a number from 1 to 128.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a VLAN interface, enter the keyword `vlan` then a number from 1 to 4094.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

2. Enable the interface.

INTERFACE mode

```
no shutdown
```

3. Configure a primary IP address and mask on the interface.

INTERFACE mode

```
ip address ip-address mask [secondary]
```

- `ip-address mask`: the IP address must be in dotted decimal format (A.B.C.D). The mask must be in slash prefix-length format (/24).
- `secondary`: add the keyword `secondary` if the IP address is the interface's backup IP address.

Example the show config Command

Example of the show ip interface Command

To view the configuration, use the `show config` command in INTERFACE mode or use the `show ip interface` command in EXEC privilege mode, as shown in the second example.

```
Dell(conf-if-te-0/16)#show conf
!  
interface TenGigabitEthernet 0/16  
  no ip address  
  shutdown  
Dell(conf-if-te-0/16)#
```

```
Dell#show ip interface tengig 0/16  
TenGigabitEthernet 0/16 is down, line protocol is down  
Internet address is not set  
IP MTU is 1500 bytes  
Directed broadcast forwarding is disabled  
Proxy ARP is enabled  
Split Horizon is enabled  
Poison Reverse is disabled  
ICMP redirects are not sent  
ICMP unreachable are not sent  
Dell#
```

Configuring Static Routes

A static route is an IP address that you manually configure and that the routing protocol does not learn, such as open shortest path first (OSPF). Often, static routes are used as backup routes in case other dynamically learned routes are unreachable.

You can enter as many static IP addresses as necessary.

To configure a static route, use the following command.

- Configure a static IP address.

CONFIGURATION mode

```
ip route ip-address mask {ip-address | interface [ip-address]} [distance]  
[permanent] [tag tag-value]
```

Use the following required and optional parameters:

- *ip-address*: enter an address in dotted decimal format (A.B.C.D).
- *mask*: enter a mask in slash prefix-length format (/X).
- *interface*: enter an interface type then the slot/port information.
- *distance*: the range is from 1 to 255. (optional)
- *permanent*: keep the static route in the routing table (if you use the *interface* option) even if you disable the interface with the route. (optional)
- *tag tag-value*: the range is from 1 to 4294967295. (optional)

Example of the show ip route static Command

To view the configured routes, use the `show ip route static` command.

```
Dell#show ip route static  
Destination Gateway Dist/Metric Last Change
```

```

-----
S 2.1.2.0/24      Direct, Nu 0          0/0      00:02:30
S 6.1.2.0/24      via 6.1.20.2, Te 5/0 1/0      00:02:30
S 6.1.2.2/32      via 6.1.20.2, Te 5/0 1/0      00:02:30
S 6.1.2.3/32      via 6.1.20.2, Te 5/0 1/0      00:02:30
S 6.1.2.4/32      via 6.1.20.2, Te 5/0 1/0      00:02:30
S 6.1.2.5/32      via 6.1.20.2, Te 5/0 1/0      00:02:30
S 6.1.2.6/32      via 6.1.20.2, Te 5/0 1/0      00:02:30
S 6.1.2.7/32      via 6.1.20.2, Te 5/0 1/0      00:02:30
S 6.1.2.8/32      via 6.1.20.2, Te 5/0 1/0      00:02:30
S 6.1.2.9/32      via 6.1.20.2, Te 5/0 1/0      00:02:30
S 6.1.2.10/32     via 6.1.20.2, Te 5/0 1/0     00:02:30
S 6.1.2.11/32     via 6.1.20.2, Te 5/0 1/0     00:02:30
S 6.1.2.12/32     via 6.1.20.2, Te 5/0 1/0     00:02:30
S 6.1.2.13/32     via 6.1.20.2, Te 5/0 1/0     00:02:30
S 6.1.2.14/32     via 6.1.20.2, Te 5/0 1/0     00:02:30
S 6.1.2.15/32     via 6.1.20.2, Te 5/0 1/0     00:02:30
S 6.1.2.16/32     via 6.1.20.2, Te 5/0 1/0     00:02:30
S 6.1.2.17/32     via 6.1.20.2, Te 5/0 1/0     00:02:30
S 11.1.1.0/24     Direct, Nu 0          0/0      00:02:30
                  Direct, Lo 0
--More--

```

The system installs a next hop that is on the directly connected subnet of current IP address on the interface (for example, if interface `tengig 0/0` is on 172.31.5.0 subnet, the system installs the static route).

The system also installs a next hop that is not on the directly connected subnet but which recursively resolves to a next hop on the interface's configured subnet. For example, if `tengig 0/0` has ip address on subnet 2.2.2.0 and if 172.31.5.43 recursively resolves to 2.2.2.0, the system installs the static route.

- When the interface goes down, the system withdraws the route.
- When the interface comes up, the system re-installs the route.
- When the recursive resolution is "broken," the system withdraws the route.
- When the recursive resolution is satisfied, the system re-installs the route.

Configure Static Routes for the Management Interface

When an IP address that a protocol uses and a static management route exists for the same prefix, the protocol route takes precedence over the static management route.

To configure a static route for the management port, use the following command.

- Assign a static route to point to the management interface or forwarding router.
CONFIGURATION mode

```

management route ip-address mask {forwarding-router-address |
ManagementEthernet slot/port}

```

Example of the `show ip route static` Command

To view the configured static routes for the management port, use the `show ip management-route` command in EXEC privilege mode.

```

Dell#show ip management-route all

Destination      Gateway          State
-----
1.1.1.0/24       172.31.1.250    Active
172.16.1.0/24    172.31.1.250    Active

```

IPv4 Path MTU Discovery Overview

The size of the packet that can be sent across each hop in the network path without being fragmented is called the path maximum transmission unit (PMTU). This value might vary for the same route between two devices, mainly over a public network, depending on the network load and speed, and it is not a consistent value. The MTU size can also be different for various types of traffic sent from one host to the same endpoint.

Path MTU discovery (PMTD) identifies the path MTU value between the sender and the receiver, and uses the determined value to transmit packets across the network. PMTD, as described in RFC 1191, denotes that the default byte size of an IP packet is 576. This packet size is called the maximum transmission unit (MTU) for IPv4 frames. PMTD operates by containing the do not fragment (DF) bit set in the IP headers of outgoing packets. When any device along the network path contains an MTU that is smaller than the size of the packet that it receives, the device drops the packet and sends an Internet Control Message Protocol (ICMP) Fragmentation Needed (Type 3, Code 4) message with its MTU value to the source or the sending device. This message enables the source to identify that the transmitted packet size must be reduced. The packet is retransmitted with a lower size than the previous value. This process is repeated in an interactive way until the MTU of the transmitted packet is lower or equal to the MTU of the receiving device for it to obtain the packet without fragmentation. If the ICMP message from the receiving device, which is sent to the originating device, contains the next-hop MTU, then the sending device lowers the packet size accordingly and resends the packet. Otherwise, the iterative method is followed until the packet can traverse without being fragmented.

PMTD is enabled by default on the switches that support this capability. To enable PMTD to function correctly, you must enter the `ip unreachable` command on a VLAN interface to enable the generation of ICMP unreachable messages. PMTD is supported on all the layer 3 VLAN interfaces. Because all of the Layer 3 interfaces are mapped to the VLAN ID of 4095 when VLAN sub-interfaces are configured on it, it is not possible to configure unique layer 3 MTU values for each of the layer 3 interfaces. If a VLAN interface contains both IPv4 and IPv6 addresses configured on it, both the IPv4 and IPv6 traffic are applied the same MTU size; you cannot specify different MTU values for IPv4 and IPv6 packets.

Using the Configured Source IP Address in ICMP Messages

ICMP error or unreachable messages are now sent with the configured IP address of the source interface instead of the front-end port IP address as the source IP address. Enable the generation of ICMP unreachable messages through the `ip unreachable` command in Interface mode. When a ping or traceroute packet from an endpoint or a device arrives at the null 0 interface configured with a static route, it is discarded. In such cases, you can configure Internet Control Message Protocol (ICMP) unreachable messages to be sent to the transmitting device.

Configuring the ICMP Source Interface

You can enable the ICMP error and unreachable messages to contain the configured IP address of the source device instead of the previous hop's IP address. This configuration helps identify the devices along the path because the DNS server maps the loopback IP address to the host name, and does not translate the IP address of every interface of the switch to the host name.

Configure the source to send the configured source interface IP address instead of using its front-end IP address in the ICMP unreachable messages and in the `traceroute` command output. Use the `ip icmp source-interface interface` or the `ipv6 icmp source-interface interface` commands in Configuration mode to enable the ICMP error messages to be sent with the source interface IP address. This functionality is supported on loopback, VLAN, port channel, and physical interfaces for IPv4 and IPv6 messages. `feature` is not supported on tunnel interfaces. ICMP error relay, PATH MTU transmission, and fragmented packets are not supported for tunnel interfaces. The traceroute utilities for IPv4 and IPv6 list the IP addresses of the devices in the hops of the path for which ICMP source interface is configured.

Configuring the Duration to Establish a TCP Connection

You can configure the amount of time for which the device must wait before it attempts to establish a TCP connection. Using this capability, you can limit the wait times for TCP connection requests. Upon responding to the initial SYN packet that requests a connection to the router for a specific service (such as SSH or BGP) with a SYN ACK, the router waits for a period of time for the ACK packet to be sent from the requesting host that will establish the TCP connection.

You can set this duration or interval for which the TCP connection waits to be established to a significantly high value to prevent the device from moving into an out-of-service condition or becoming unresponsive during a SYN flood attack that occurs on the device. You can set the wait time to be 10 seconds or lower. If the device does not contain any BGP connections with the BGP neighbors across WAN links, you must set this interval to a higher value, depending on the complexity of your network and the configuration attributes.

To configure the duration for which the device waits for the ACK packet to be sent from the requesting host to establish the TCP connection, perform the following steps:

1. Define the wait duration in seconds for the TCP connection to be established.

CONFIGURATION mode

```
Dell(conf)#ip tcp reduced-syn-ack-wait <9-75>
```

You can use the `no ip tcp reduced-syn-ack-wait` command to restore the default behavior, which causes the wait period to be set as 8 seconds.

2. View the interval that you configured for the device to wait before the TCP connection is attempted to be established.

EXEC mode

```
Dell>show ip tcp reduced-syn-ack-wait
```

Enabling Directed Broadcast

By default, the system drops directed broadcast packets destined for an interface. This default setting provides some protection against denial of service (DoS) attacks.

To enable the system to receive directed broadcasts, use the following command.

- Enable directed broadcast.
INTERFACE mode

```
ip directed-broadcast
```

To view the configuration, use the `show config` command in INTERFACE mode.

Resolution of Host Names

Domain name service (DNS) maps host names to IP addresses. This feature simplifies such commands as Telnet and FTP by allowing you to enter a name instead of an IP address.

Dynamic resolution of host names is disabled by default. Unless you enable the feature, the system resolves only host names entered into the host table with the `ip host` command.

The following sections describe DNS and the resolution of host names.

- [Enabling Dynamic Resolution of Host Names](#)
- [Specifying the Local System Domain and a List of Domains](#)
- [Configuring DNS with Traceroute](#)

Enabling Dynamic Resolution of Host Names

By default, dynamic resolution of host names (DNS) is disabled.

To enable DNS, use the following commands.

- Enable dynamic resolution of host names.
CONFIGURATION mode

```
ip domain-lookup
```
- Specify up to six name servers.
CONFIGURATION mode

```
ip name-server ip-address [ip-address2 ... ip-address6]
```

The order you entered the servers determines the order of their use.

Example of the `show hosts` Command

To view current bindings, use the `show hosts` command.

```
Dell>show host
Default domain is force10networks.com
Name/address lookup uses domain service
```

```

Name servers are not set
Host      Flags TTL      Type Address
-----
ks        (perm, OK) - IP 2.2.2.2
patch1    (perm, OK) - IP 192.68.69.2
tomm-3    (perm, OK) - IP 192.68.99.2
gxr       (perm, OK) - IP 192.71.18.2
f00-3     (perm, OK) - IP 192.71.23.1
Dell>

```

To view the current configuration, use the `show running-config resolve` command.

Specifying the Local System Domain and a List of Domains

If you enter a partial domain, the system can search different domains to finish or fully qualify that partial domain.

A fully qualified domain name (FQDN) is any name that is terminated with a period/dot. The Dell Networking OS searches the host table first to resolve the partial domain. The host table contains both statically configured and dynamically learnt host and IP addresses. If the system cannot resolve the domain, it tries the domain name assigned to the local system. If that does not resolve the partial domain, the system searches the list of domains configured.

To configure a domain name or a list of domain names, use the following commands.

- Enter up to 63 characters to configure one domain name.

CONFIGURATION mode

```
ip domain-name name
```

- Enter up to 63 characters to configure names to complete unqualified host names.

CONFIGURATION mode

```
ip domain-list name
```

Configure this command up to six times to specify a list of possible domain names. The Dell Networking OS searches the domain names in the order they were configured until a match is found or the list is exhausted.

Configuring DNS with Traceroute

To configure your switch to perform DNS with traceroute, use the following commands.

- Enable dynamic resolution of host names.

CONFIGURATION mode

```
ip domain-lookup
```

- Specify up to six name servers.

CONFIGURATION mode

```
ip name-server ip-address [ip-address2 ... ip-address6]
```

The order you entered the servers determines the order of their use.

- When you enter the `traceroute` command without specifying an IP address (Extended Traceroute), you are prompted for a target and source IP address, timeout in seconds (default is **5**),

a probe count (default is **3**), minimum TTL (default is **1**), maximum TTL (default is **30**), and port number (default is **33434**).

CONFIGURATION mode

```
traceroute [host | ip-address]
```

To keep the default setting for these parameters, press the ENTER key.

Example of the `traceroute` Command

The following text is example output of DNS using the `traceroute` command.

```
Dell#traceroute www.force10networks.com
```

```
Translating "www.force10networks.com"...domain server (10.11.0.1) [OK]
Type Ctrl-C to abort.
```

```
-----
-----
Tracing the route to www.force10networks.com (10.11.84.18), 30 hops max, 40
byte packets
-----
-----
```

TTL	Hostname	Probe1	Probe2	Probe3
1	10.11.199.190	001.000 ms	001.000 ms	002.000 ms
2	gwegress-sjc-02.force10networks.com (10.11.30.126)	005.000 ms	001.000 ms	001.000 ms
3	fw-sjc-01.force10networks.com (10.11.127.254)	000.000 ms	000.000 ms	000.000 ms
4	www.force10networks.com (10.11.84.18)	000.000 ms	000.000 ms	000.000 ms

```
Dell#
```

ARP

The Dell Networking OS uses two forms of address resolution: address resolution protocol (ARP) and Proxy ARP.

ARP runs over Ethernet and enables endstations to learn the MAC addresses of neighbors on an IP network. Over time, the system creates a forwarding table mapping the MAC addresses to their corresponding IP address. This table is called the ARP Cache and dynamically learned addresses are removed after a defined period of time.

For more information about ARP, refer to RFC 826, *An Ethernet Address Resolution Protocol*.

In the Dell Networking OS, Proxy ARP enables hosts with knowledge of the network to accept and forward packets from hosts that contain no knowledge of the network. Proxy ARP makes it possible for hosts to be ignorant of the network, including subnetting.

For more information about Proxy ARP, refer to RFC 925, *Multi-LAN Address Resolution*, and RFC 1027, *Using ARP to Implement Transparent Subnet Gateways*.

Configuration Tasks for ARP

For a complete listing of all ARP-related commands, refer to the *Dell Networking OS Command Line Reference Guide*.

Configuration tasks for ARP include:

- [Configuring Static ARP Entries](#) (optional)
- [Enabling Proxy ARP](#) (optional)
- [Clearing ARP Cache](#) (optional)
- [ARP Learning via Gratuitous ARP](#)
- [ARP Learning via ARP Request](#)
- [Configuring ARP Retries](#)

Configuring Static ARP Entries

ARP dynamically maps the MAC and IP addresses, and while most network host support dynamic mapping, you can configure an ARP entry (called a static ARP) for the ARP cache.

To configure a static ARP entry, use the following command.

- Configure an IP address and MAC address mapping for an interface.

CONFIGURATION mode

```
arp ip-address mac-address interface
```

- *ip-address*: IP address in dotted decimal format (A.B.C.D).
- *mac-address*: MAC address in nnnn.nnnn.nnnn format.
- *interface*: enter the interface type slot/port information.

Example of the `show arp` Command

These entries do not age and can only be removed manually. To remove a static ARP entry, use the `no arp ip-address` command.

To view the static entries in the ARP cache, use the `show arp static` command in EXEC privilege mode.

```
Dell#show arp
```

Protocol	Address	Age (min)	Hardware Address	Interface	VLAN	CPU

Internet	10.11.68.14	94	00:01:e9:45:00:03	Ma 0/0	-	CP
Internet	10.11.209.254	0	00:01:e9:45:00:03	Ma 0/0	-	CP

```
Dell#
```

Enabling Proxy ARP

By default, Proxy ARP is enabled. To disable Proxy ARP, use the `no proxy-arp` command in the interface mode.

To re-enable Proxy ARP, use the following command.

- Re-enable Proxy ARP.

INTERFACE mode

```
ip proxy-arp
```

To view if Proxy ARP is enabled on the interface, use the `show config` command in INTERFACE mode. If it is not listed in the `show config` command output, it is enabled. Only non-default information is displayed in the `show config` command output.


Clearing ARP Cache

To clear the ARP cache of dynamically learnt ARP information, use the following command.

- Clear the ARP caches for all interfaces or for a specific interface by entering the following information. EXEC privilege

```
clear arp-cache [interface | ip ip-address] [no-refresh]
```

- `ip ip-address` (OPTIONAL): enter the keyword `ip` then the IP address of the ARP entry you wish to clear.
- `no-refresh` (OPTIONAL): enter the keywords `no-refresh` to delete the ARP entry from CAM. Or to specify which dynamic ARP entries you want to delete, use this option with `interface` or `ip ip-address`.
- For a port channel interface, enter the keywords `port-channel` then a number from 1 to 128.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a VLAN interface, enter the keyword `vlan` then a number between 1 and 4094.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

 **NOTE:** Transit traffic may not be forwarded during the period when deleted ARP entries are resolved again and re-installed in CAM. Use this option with extreme caution.

ARP Learning via Gratuitous ARP

Gratuitous ARP can mean an ARP request or reply.

In the context of ARP learning via gratuitous ARP on the system, the gratuitous ARP is a request. A gratuitous ARP request is an ARP request that is not needed according to the ARP specification, but one that hosts may send to. Gratuitous ARP can:

- detect IP address conflicts
- inform switches of their presence on a port so that packets can be forwarded
- update the ARP table of other nodes on the network in case of an address change

In the request, the host uses its own IP address in the Sender Protocol Address and Target Protocol Address fields. When a gratuitous ARP is received, the system installs an ARP entry on the CPU.

To enable ARP learning via gratuitous ARP, use the `arp learn-enable` command in CONFIGURATION mode.

ARP Learning via ARP Request

In the Dell Networking OS versions prior to 8.3.1.0, the system learns via ARP requests only if the target IP specified in the packet matches the IP address of the receiving router interface. This is the case when a host is attempting to resolve the gateway address.

If the target IP does not match the incoming interface, the packet is dropped. If there is an existing entry for the requesting host, it is updated.




Figure 48. ARP Learning via ARP Request

Beginning with the Dell Networking OS version 8.3.1.0, when you enable ARP learning via gratuitous ARP, the system installs a new ARP entry, or updates an existing entry for all received ARP requests.

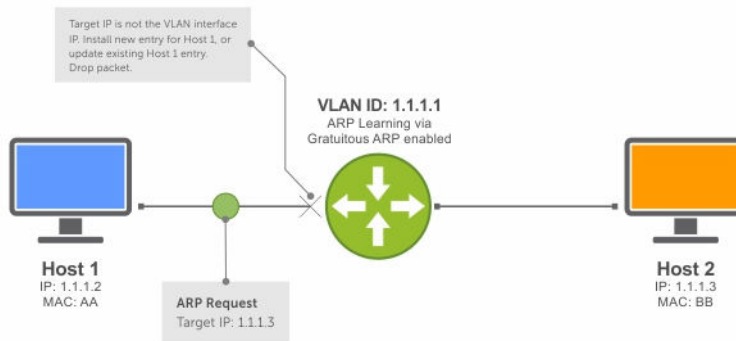


Figure 49. ARP Learning via ARP Request with ARP Learning via Gratuitous ARP Enabled

Whether you enable or disable ARP learning via gratuitous ARP, the system does not look up the target IP. It only updates the ARP entry for the Layer 3 interface with the source IP of the request.

Configuring ARP Retries

In the Dell Networking OS versions prior to 8.3.1.0, the number of ARP retries is set to five and is not configurable. After five retries, the system backs off for 20 seconds before it sends a new request. Beginning with the Dell Networking OS version 8.3.1.0, the number of ARP retries is configurable. The default backoff interval remains at 20 seconds. On the MXL switch platform, using the Dell Networking OS version 8.3.8.0 and later, the time between ARP re-send is configurable. This timer is an exponential backoff timer. Over the specified period, the time between ARP requests increases. This reduces the potential for the system to slow down while waiting for a multitude of ARP responses.

To set and display ARP retries, use the following commands.

- Set the number of ARP retries.

CONFIGURATION mode

```
arp retries number
```

The default is **5**.

The range is from 1 to 20.

- Set the exponential timer for resending unresolved ARPs.

CONFIGURATION mode

```
arp backoff-time
```

The default is **30**.

The range is from 1 to 3600.

- Display all ARP entries learned via gratuitous ARP.

EXEC Privilege mode

```
show arp retries
```

ICMP

For diagnostics, the internet control message protocol (ICMP) provides routing information to end stations by choosing the best route (ICMP redirect messages) or determining if a router is reachable (ICMP Echo or Echo Reply).

ICMP error messages inform the router of problems in a particular packet. These messages are sent only on unicast traffic.

Configuration Tasks for ICMP

The following lists the configuration tasks for ICMP.

- [Enabling ICMP Unreachable Messages](#)

For a complete listing of all commands related to ICMP, refer to the *Dell Networking OS Command Line Reference Guide*.

Enabling ICMP Unreachable Messages

By default, ICMP unreachable messages are disabled.

When enabled, ICMP unreachable messages are created and sent out all interfaces.

To disable and re-enable ICMP unreachable messages, use the following commands.

- Set the system to create and send ICMP unreachable messages on the interface.

```
INTERFACE mode
```

```
ip unreachable
```

To view if ICMP unreachable messages are sent on the interface, use the `show config` command in INTERFACE mode. If it is not listed in the `show config` command output, it is enabled. Only non-default information is displayed in the `show config` command output.

UDP Helper

User datagram protocol (UDP) helper allows you to direct the forwarding IP/UDP broadcast traffic by creating special broadcast addresses and rewriting the destination IP address of packets to match those addresses.

Configure UDP Helper

Configuring the system to direct UDP broadcast is a one-step process:

1. Enable UDP helper and specify the UDP ports for which traffic is forwarded. Refer to [Enabling UDP Helper](#).

Important Points to Remember

- The existing `ip directed broadcast` command is rendered meaningless if you enable UDP helper on the same interface.
- The broadcast traffic rate should not exceed 200 packets per second when you enable UDP helper.
- You may specify a maximum of 16 UDP ports.
- UDP helper is compatible with IP helper (`ip helper-address`):
 - UDP broadcast traffic with port number 67 or 68 are unicast to the dynamic host configuration protocol (DHCP) server per the `ip helper-address` configuration whether or not the UDP port list contains those ports.
 - If the UDP port list contains ports 67 or 68, UDP broadcast traffic is forwarded on those ports.

Enabling UDP Helper

To enable UDP helper, use the following command.

- Enable UDP helper.

```
ip udp-helper udp-ports
```

Example of Enabling UDP Helper

Example of the `show ip udp-helper` Command

```
Dell(conf-if-te-1/1)#ip udp-helper udp-port 1000
Dell(conf-if-te-1/1)#show config
!
interface TenGigabitEthernet 1/1
 ip address 2.1.1.1/24
 ip udp-helper udp-port 1000
 no shutdown
```

To view the interfaces and ports on which you enabled UDP helper, use the `show ip udp-helper` command from EXEC Privilege mode.

```
Dell#show ip udp-helper
-----
Port          UDP port list
-----
TenGig 1/1    1000
```

Configurations Using UDP Helper

When you enable UDP helper and the destination IP address of an incoming packet is a broadcast address, the system suppresses the destination address of the packet.

The following sections describe various configurations that employ UDP helper to direct broadcasts.

- [UDP Helper with Broadcast-All Addresses](#)
- [UDP Helper with Subnet Broadcast Addresses](#)
- [UDP Helper with Configured Broadcast Addresses](#)
- [UDP Helper with No Configured Broadcast Addresses](#)

UDP Helper with Broadcast-All Addresses

When the destination IP address of an incoming packet is the IP broadcast address, the system rewrites the address to match the configured broadcast address.

In the following illustration:

1. Packet 1 is dropped at ingress if you did not configure UDP helper address.
2. If you enable UDP helper (using the `ip udp-helper udp-port` command), and the UDP destination port of the packet matches the UDP port configured, the system changes the destination address to the configured broadcast 1.1.255.255 and routes the packet to VLANs 100 and 101. If you do not configure an IP broadcast address (using the `ip udp-broadcast-address` command) on VLANs 100 or 101, the packet is forwarded using the original destination IP address 255.255.255.255.

Packet 2, sent from a host on VLAN 101 has a broadcast MAC address and IP address. In this case:

1. It is flooded on VLAN 101 without changing the destination address because the forwarding process is Layer 2.
2. If you enabled UDP helper, the system changes the destination IP address to the configured broadcast address 1.1.255.255 and forwards the packet to VLAN 100.

- Packet 2 is also forwarded to the ingress interface with an unchanged destination address because it does not have broadcast address configured.

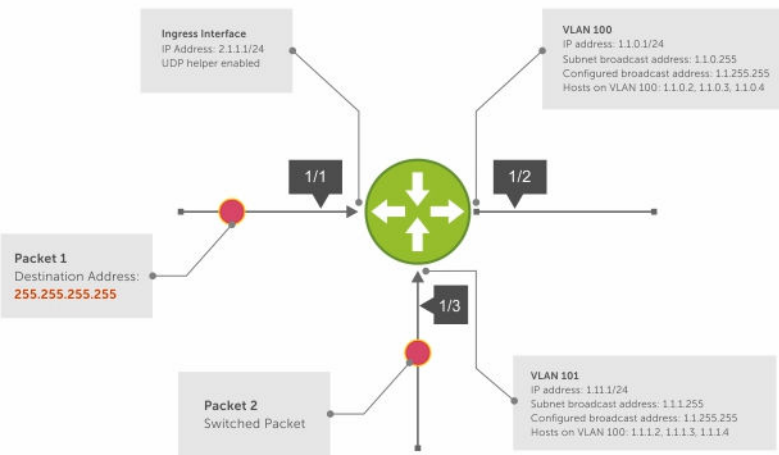


Figure 50. UDP Helper with Broadcast-All Addresses

UDP Helper with Subnet Broadcast Addresses

When the destination IP address of an incoming packet matches the subnet broadcast address of any interface, the system changes the address to the configured broadcast address and sends it to matching interface.

In the following illustration, Packet 1 has the destination IP address 1.1.1.255, which matches the subnet broadcast address of VLAN 101. If you configured UDP helper and the packet matches the specified UDP port, the system changes the address to the configured IP broadcast address and floods the packet on VLAN 101.

Packet 2 is sent from the host on VLAN 101. It has a broadcast MAC address and a destination IP address of 1.1.1.255. In this case, it is flooded on VLAN 101 in its original condition as the forwarding process is Layer 2.




Figure 51. UDP Helper with Subnet Broadcast Addresses

UDP Helper with Configured Broadcast Addresses

Incoming packets with a destination IP address matching the configured broadcast address of any interface are forwarded to the matching interfaces.

In the following illustration, Packet 1 has a destination IP address that matches the configured broadcast address of VLAN 100 and 101. If you enabled UDP helper and the UDP port number matches, the packet is flooded on both VLANs with an unchanged destination address.

Packet 2 is sent from a host on VLAN 101. It has broadcast MAC address and a destination IP address that matches the configured broadcast address on VLAN 101. In this case, Packet 2 is flooded on VLAN 101 with the destination address unchanged because the forwarding process is Layer 2. If you enabled UDP helper, the packet is flooded on VLAN 100 as well.

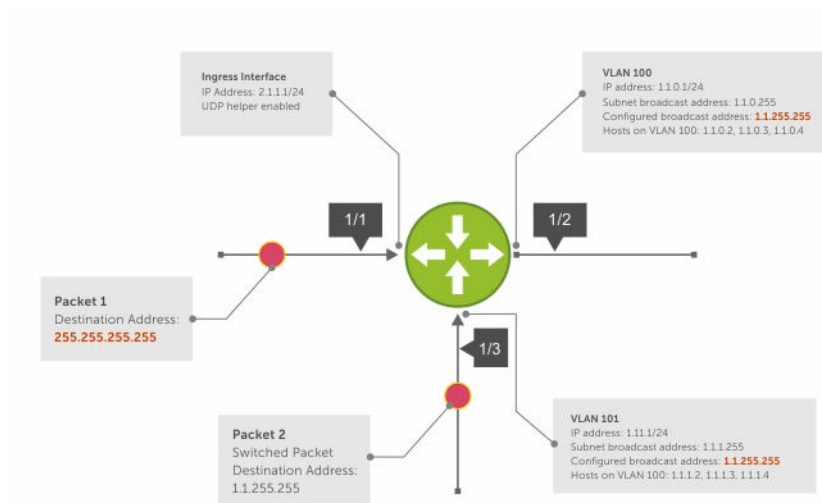


Figure 52. UDP Helper with Configured Broadcast Addresses

UDP Helper with No Configured Broadcast Addresses

The following describes UDP helper with no broadcast addresses configured.

- If the incoming packet has a broadcast destination IP address, the unaltered packet is routed to all Layer 3 interfaces.
- If the Incoming packet has a destination IP address that matches the subnet broadcast address of any interface, the unaltered packet is routed to the matching interfaces.

Troubleshooting UDP Helper

To display debugging information for troubleshooting, use the `debug ip udp-helper` command.

Example of the `debug ip udp-helper` Command

```
Dell(conf)# debug ip udp-helper
01:20:22: Pkt rcvd on TenGig 5/0 with IP DA (0xffffffff) will be sent on TenGig
```

```
5/1 TenGig 5/  
2 Vlan 3  
01:44:54: Pkt rcvd on TenGig 7/0 is handed over for DHCP processing.
```


When using the IP helper and UDP helper on the same interface, use the `debug ip dhcp` command.

Example Output from the `debug ip dhcp` Command

```
Packet 0.0.0.0:68 -> 255.255.255.255:67 TTL 128  
  
2005-11-05 11:59:35 %RELAY-I-PACKET, BOOTP REQUEST (Unicast) received at  
interface  
172.21.50.193 BOOTP Request, XID = 0x9265f901, secs = 0 hwaddr = 00:02:2D:8D:  
46:DC, giaddr =  
0.0.0.0, hops = 2  
  
2005-11-05 11:59:35 %RELAY-I-BOOTREQUEST, Forwarded BOOTREQUEST for 00:02:2D:8D:  
46:DC to  
137.138.17.6  
  
2005-11-05 11:59:36 %RELAY-I-PACKET, BOOTP REPLY (Unicast) received at interface  
194.12.129.98 BOOTP Reply, XID = 0x9265f901, secs = 0 hwaddr = 00:02:2D:8D:  
46:DC, giaddr =  
172.21.50.193, hops = 2  
  
2005-07-05 11:59:36 %RELAY-I-BOOTREPLY, Forwarded BOOTREPLY for 00:02:2D:8D:  
46:DC to  
128.141.128.90 Packet 0.0.0.0:68 -> 255.255.255.255:67 TTL 128
```


IPv6 Addressing

Internet protocol version 6 (IPv6) is supported on the MXL switch platform.

 **NOTE:** The IPv6 basic commands are supported on all platforms. However, not all features are supported on all platforms, nor for all releases. To determine the Dell Networking OS version supporting which features and platforms, refer to [Implementing IPv6 with the Dell Networking OS](#).

IPv6 is the successor to IPv4. Due to the rapid growth in internet users and IP addresses, IPv4 is reaching its maximum usage. IPv6 will eventually replace IPv4 usage to allow for the constant expansion.

This chapter provides a brief description of the differences between IPv4 and IPv6, and the Dell Networking support of IPv6. This chapter is not intended to be a comprehensive description of IPv6.

 **NOTE:** Even though Dell Networking OS listens to all ports, you can only use the ports starting from 1024 for IPv6 traffic. Ports from 0 to 1023 are reserved for internal use and you cannot use them for IPv6 traffic.

Protocol Overview

IPv6 is an evolution of IPv4. IPv6 is generally installed as an upgrade in devices and operating systems. Most new devices and operating systems support both IPv4 and IPv6.

Some key changes in IPv6 are:

- Extended address space
- Stateless autoconfiguration
- Header format simplification
- Improved support for options and extensions

Extended Address Space

The address format is extended from 32 bits to 128 bits. This not only provides room for all anticipated needs, it allows for the use of a hierarchical address space structure to optimize global addressing.


Stateless Autoconfiguration


When a booting device comes up in IPv6 and asks for its network prefix, the device can get the prefix (or prefixes) from an IPv6 router on its link. It can then autoconfigure one or more global IPv6 addresses by using either the MAC address or a private random number to build its unique IPv6 address.

Stateless autoconfiguration uses three mechanisms for IPv6 address configuration:


- **Prefix Advertisement** — Routers use “Router Advertisement” messages to announce the network prefix. Hosts then use their interface-identifier MAC address to generate their own valid IPv6 address.

- **Duplicate Address Detection (DAD)** — Before configuring its IPv6 address, an IPv6 host node device checks whether that address is used anywhere on the network using this mechanism.
- **Prefix Renumbering** — Useful in transparent renumbering of hosts in the network when an organization changes its service provider.

 **NOTE:** As an alternative to stateless autoconfiguration, network hosts can obtain their IPv6 addresses using the dynamic host control protocol (DHCP) servers via stateful auto-configuration.

 **NOTE:** The Dell Networking OS provides the flexibility to add prefixes on Router Advertisements (RA) to advertise responses to Router Solicitations (RS). By default, RA response messages are sent when an RS message is received.

The Dell Networking OS manipulation of IPv6 stateless autoconfiguration supports the router side only. Neighbor discovery (ND) messages are advertised so the neighbor can use this information to auto-configure its address. However, received ND messages are not used to create an IPv6 address.

 **NOTE:** Inconsistencies in router advertisement values between routers are logged per RFC 4861. The values checked for consistency include:

- Cur Hop limit
- M and O flags
- Reachable time
- Retrans timer
- MTU options
- Preferred and valid lifetime values for the same prefix

Only management ports support stateless auto-configuration as a host.

The router redirect functionality in the neighbor discovery protocol (NDP) is similar to IPv4 router redirect messages. NDP uses ICMPv6 redirect messages (Type 137) to inform nodes that a better router exists on the link.

IPv6 Headers

The IPv6 header has a fixed length of 40 bytes. This fixed length provides 16 bytes each for source and destination information and 8 bytes for general header information.

The IPv6 header includes the following fields:

- [Version \(4 bits\)](#)
- [Traffic Class \(8 bits\)](#)
- [Flow Label \(20 bits\)](#)
- [Payload Length \(16 bits\)](#)
- [Next Header \(8 bits\)](#)
- [Hop Limit \(8 bits\)](#)
- [Source Address \(128 bits\)](#)
- [Destination Address \(128 bits\)](#)

IPv6 provides for extension headers. Extension headers are used only if necessary. There can be no extension headers, one extension header or more than one extension header in an IPv6 packet. Extension headers are defined in the Next Header field of the preceding IPv6 header.

IPv6 Header Fields

The 40 bytes of the IPv6 header are ordered, as shown in the following illustration.



Figure 53. IPv6 Header Fields

Version (4 bits)

The Version field always contains the number 6, referring to the packet's IP version.


Traffic Class (8 bits)

The Traffic Class field deals with any data that needs special handling. These bits define the packet priority and are defined by the packet Source. Sending and forwarding routers use this field to identify different IPv6 classes and priorities. Routers understand the priority settings and handle them appropriately during conditions of congestion.

Flow Label (20 bits)

The Flow Label field identifies packets requiring special treatment in order to manage real-time data traffic.

The sending router can label sequences of IPv6 packets so that forwarding routers can process packets within the same flow without needing to reprocess each packet's header separately.

 **NOTE:** All packets in the flow must have the same source and destination addresses.

Payload Length (16 bits)

The Payload Length field specifies the packet payload. This is the length of the data following the IPv6 header. IPv6 Payload Length only includes the data following the header, not the header itself.

The Payload Length limit of 2 bytes requires that the maximum packet payload be 64 KB. However, the Jumbogram option type Extension header supports larger packet sizes when required.

Next Header (8 bits)

The Next Header field identifies the next header's type. If an Extension header is used, this field contains the type of Extension header (as shown in the following table). If the next header is a transmission control protocol (TCP) or user datagram protocol (UDP) header, the value in this field is the same as for IPv4. The Extension header is located between the IP header and the TCP or UDP header.

The following lists the Next Header field values.

Value	Description
0	Hop-by-Hop option header
4	IPv4
6	TCP
8	Exterior Gateway Protocol (EGP)
41	IPv6
43	Routing header
44	Fragmentation header
50	Encrypted Security
51	Authentication header
59	No Next Header
60	Destinations option header

 **NOTE:** This table is not a comprehensive list of Next Header field values. For a complete and current listing, refer to the Internet Assigned Numbers Authority (IANA) web page.

Hop Limit (8 bits)

The Hop Limit field shows the number of hops remaining for packet processing. In IPv4, this is known as the Time to Live (TTL) field and uses seconds rather than hops.

Each time the packet moves through a forwarding router, this field decrements by 1. If a router receives a packet with a Hop Limit of 1, it decrements it to 0 (zero). The router discards the packet and sends an ICMPv6 message back to the sending router indicating that the Hop Limit was exceeded in transit.

Source Address (128 bits)

The Source Address field contains the IPv6 address for the packet originator.

Destination Address (128 bits)

The Destination Address field contains the intended recipient's IPv6 address. This can be either the ultimate destination or the address of the next hop router.

Extension Header Fields

Extension headers are used only when necessary. Due to the streamlined nature of the IPv6 header, adding extension headers do not severely impact performance. Each Extension headers's lengths vary, but they are always a multiple of 8 bytes.

Each extension header is identified by the Next Header field in the IPv6 header that precedes it. Extension headers are viewed only by the destination router identified in the Destination Address field. If the Destination Address is a multicast address, the Extension headers are examined by all the routers in that multicast group.

However, if the Destination Address is a Hop-by-Hop options header, the Extension header is examined by every forwarding router along the packet's route. The Hop-by-Hop options header must immediately follow the IPv6 header, and is noted by the value 0 (zero) in the Next Header field.

Extension headers are processed in the order in which they appear in the packet header.

Hop-by-Hop Options Header

The Hop-by-Hop options header contains information that is examined by every router along the packet's path. It follows the IPv6 header and is designated by the Next Header value 0 (zero).

When a Hop-by-Hop Options header is not included, the router knows that it does not have to process any router specific information and immediately processes the packet to its final destination.

When a Hop-by-Hop Options header is present, the router only needs this extension header and does not need to take the time to view further into the packet.

The Hop-by-Hop Options header contains:

- Next Header (1 byte)

This field identifies the type of header following the Hop-by-Hop Options header and uses the same values.

- Header Extension Length (1 byte)

This field identifies the length of the Hop-by-Hop Options header in 8-byte units, but does not include the first 8 bytes. Consequently, if the header is less than 8 bytes, the value is 0 (zero).

- Options (size varies)

This field can contain one or more options. The first byte of the field identifies the Option type, and directs the router how to handle the option.

00	Skip and continue processing.
01	Discard the packet.
10	Discard the packet and send an ICMP Parameter Problem Code 2 message to the packet's Source IP Address identifying the unknown option type.
11	Discard the packet and send an ICMP Parameter Problem, Code 2 message to the packet's Source IP Address only if the Destination IP Address is not a multicast address.

The second byte contains the Option Data Length.

The third byte specifies whether the information can change en route to the destination. The value is 1 if it can change; the value is 0 if it cannot change.

Addressing

IPv6 addresses are normally written as eight groups of four hexadecimal digits, where each group is separated by a colon (:).

For example, 2001:0db8:0000:0000:0000:0000:1428:57ab is a valid IPv6 address. If one or more four-digit group(s) is 0000, the zeros may be omitted and replaced with two colons (::). For example, 2001:0db8:0000:0000:0000:0000:1428:57ab can be shortened to 2001:0db8::1428:57ab. Only one set of double colons is supported in a single address. Any number of consecutive 0000 groups may be reduced to two colons, as long as there is only one double colon used in an address. Leading and/or trailing zeros in a group can also be omitted (as in ::1 for localhost, 1:: for network addresses and :: for unspecified addresses).

All the addresses in the following list are all valid and equivalent.

- 2001:0db8:0000:0000:0000:0000:1428:57ab
- 2001:0db8:0000:0000:0000::1428:57ab
- 2001:0db8:0:0:0:0:1428:57ab
- 2001:0db8:0:0::1428:57ab
- 2001:0db8::1428:57ab
- 2001:db8::1428:57ab

IPv6 networks are written using classless inter-domain routing (CIDR) notation. An IPv6 network (or subnet) is a contiguous group of IPv6 addresses the size of which must be a power of two; the initial bits of addresses, which are identical for all hosts in the network, are called the network's prefix.

A network is denoted by the first address in the network and the size in bits of the prefix (in decimal), separated with a slash. Because a single host is seen as a network with a 128-bit prefix, host addresses may be written with a following /128.

For example, 2001:0db8:1234::/48 stands for the network with addresses 2001:0db8:1234:0000:0000:0000:0000:0000 through 2001:0db8:1234:ffff:ffff:ffff:ffff:ffff.

Link-local Addresses

Link-local addresses, starting with fe80:, are assigned only in the local link area.

The addresses are generated usually automatically by the operating system's IP layer for each network interface. This provides instant automatic network connectivity for any IPv6 host and means that if several hosts connect to a common hub or switch, they have an instant communication path via their link-local IPv6 address.

Link-local addresses cannot be routed to the public Internet.

Static and Dynamic Addressing

Static IPv6 addresses are manually assigned to a computer by an administrator.

Dynamic IPv6 addresses are assigned either randomly or by a server using dynamic host configuration protocol (DHCP). Even though IPv6 addresses assigned using DHCP may stay the same for long periods of time, they can change. In some cases, a network administrator may implement dynamically assigned static IPv6 addresses. In this case, a DHCP server is used, but it is specifically configured to always assign the same IPv6 address to a particular computer, and never to assign that IP address to another computer. This allows static IPv6 addresses to be configured in one place, without having to specifically configure each computer on the network in a different way.

In IPv6, every interface, whether using static or dynamic address assignments, also receives a local-link address automatically in the fe80::/64 subnet.

Implementing IPv6 with the Dell Networking OS

The Dell Networking OS supports both IPv4 and IPv6 and both may be used simultaneously in your system.

The following table lists the Dell Networking OS version in which an IPv6 feature became available for each platform. The sections following the table give greater detail about the feature.

Feature and Functionality	Dell Networking OS Release Introduction	Documentation and Chapter Location
MXL		
Basic IPv6 Commands	9.2(0.0)	IPv6 Basic Commands in the <i>Dell Networking OS Command Line Interface Reference Guide</i> .
IPv6 Basic Addressing		
IPv6 address types: Unicast	9.2(0.0)	Extended Address Space in this chapter
IPv6 neighbor discovery	9.2(0.0)	IPv6 Neighbor Discovery in this chapter

Feature and Functionality	Dell Networking OS Release Introduction	Documentation and Chapter Location
MXL		
IPv6 stateless autoconfiguration	9.2(0.0)	Stateless Autoconfiguration in this chapter
IPv6 MTU path discovery	9.2(0.0)	Path MTU Discovery in this chapter
IPv6 ICMPv6	9.2(0.0)	ICMPv6 in this chapter
IPv6 ping	9.2(0.0)	ICMPv6 in this chapter
IPv6 traceroute	9.2(0.0)	ICMPv6 in this chapter
IPv6 Routing		
Static routing	9.2(0.0)	Assigning a Static IPv6 Route in this chapter
Route redistribution	9.2(0.0)	OSPF, IS-IS, and IPv6 BGP chapters in the <i>Dell Networking OS Command Line Reference Guide</i> .
Multiprotocol BGP extensions for IPv6	9.2(0.0)	IPv6 BGP in the <i>Dell Networking OS Command Line Reference Guide</i> .
IPv6 BGP MD5 Authentication	9.2(0.0)	IPv6 BGP in the <i>Dell Networking OS Command Line Reference Guide</i> .
IS-IS for IPv6	9.2(0.0)	Intermediate System to Intermediate System (IS-IS) IPv6 IS-IS in the <i>Dell Networking OS Command Line Reference Guide</i> .
IS-IS for IPv6 support for redistribution	9.2(0.0)	Intermediate System to Intermediate System (IS-IS) IPv6 IS-IS in the <i>Dell Networking OS Command Line Reference Guide</i> .
ISIS for IPv6 support for distribute lists and administrative distance	9.2(0.0)	Intermediate System to Intermediate System (IS-IS) IPv6 IS-IS in the <i>Dell Networking OS Command Line Reference Guide</i> .
OSPF for IPv6 (OSPFv3)	9.2(0.0)	OSPFv3 in the <i>Dell Networking OS Command Line Reference Guide</i> .
Equal Cost Multipath for IPv6	9.2(0.0)	

Feature and Functionality	Dell Networking OS Release Introduction	Documentation and Chapter Location
MXL		
IPv6 Services and Management		
Telnet client over IPv6 (outbound Telnet)	9.2(0.0)	Configuring Telnet with IPv6 in this chapter Control and Monitoring in the <i>Dell Networking OS Command Line Reference Guide</i> .
Telnet server over IPv6 (inbound Telnet)	9.2(0.0)	Configuring Telnet with IPv6 in this chapter Control and Monitoring in the <i>Dell Networking OS Command Line Reference Guide</i> .
Secure Shell (SSH) client support over IPv6 (outbound SSH) Layer 3 only	9.2(0.0)	Secure Shell (SSH) Over an IPv6 Transport in this chapter
Secure Shell (SSH) server support over IPv6 (inbound SSH) Layer 3 only	9.2(0.0)	Secure Shell (SSH) Over an IPv6 Transport in this chapter
IPv6 Access Control Lists	9.2(0.0)	IPv6 Access Control Lists in the <i>Dell Networking OS Command Line Reference Guide</i> .
IPv6 Multicast		
PIM-SM for IPv6	N/A	IPv6 Multicast in this chapter IPv6 PIM in the <i>Dell Networking OS Command Line Reference Guide</i> .
PIM-SSM for IPv6	N/A	IPv6 Multicast in this chapter IPv6 PIM in the <i>Dell Networking OS Command Line Reference Guide</i> .
MLDv1/v2	N/A	IPv6 Multicast in this chapter Multicast IPv6 in the <i>Dell Networking OS Command Line Reference Guide</i> .
MLDv1 Snooping	N/A	IPv6 Multicast in this chapter Multicast IPv6 in the <i>Dell Networking OS Command Line Reference Guide</i> .
MLDv2 Snooping	N/A	IPv6 Multicast in this chapter

Feature and Functionality	Dell Networking OS Release Introduction	Documentation and Chapter Location
	MXL	

Multicast IPv6 in the *Dell Networking OS Command Line Reference Guide*.

IPv6 QoS

trust DSCP values	N/A	IPv6 Multicast in this chapter
-------------------	-----	--

ICMPv6

ICMPv6 is supported on the MXL switch platform.

ICMP for IPv6 combines the roles of ICMP, IGMP and ARP in IPv4. Similar to IPv4, it provides functions for reporting delivery and forwarding errors, and provides a simple echo service for troubleshooting. The Dell Networking OS implementation of ICMPv6 is based on RFC 4443.

Generally, ICMPv6 uses two message types:

- Error reporting messages indicate when the forwarding or delivery of the packet failed at the destination or intermediate node. These messages include Destination Unreachable, Packet Too Big, Time Exceeded and Parameter Problem messages.
- Informational messages provide diagnostic functions and additional host functions, such as Neighbor Discovery and Multicast Listener Discovery. These messages also include Echo Request and Echo Reply messages.

The `ping` and `traceroute` commands extend to support IPv6 addresses. These commands use ICMPv6 Type-2 messages.

Path MTU Discovery

IPv6 path maximum transmission unit (MTU) discovery is supported on the MXL switch platform.

Path MTU, in accordance with RFC 1981, defines the largest packet size that can traverse a transmission path without suffering fragmentation. Path MTU for IPv6 uses ICMPv6 Type-2 messages to discover the largest MTU along the path from source to destination and avoid the need to fragment the packet.

The recommended MTU for IPv6 is 1280. Greater MTU settings increase processing efficiency because each packet carries more data while protocol overheads (for example, headers) or underlying per-packet delays remain fixed.

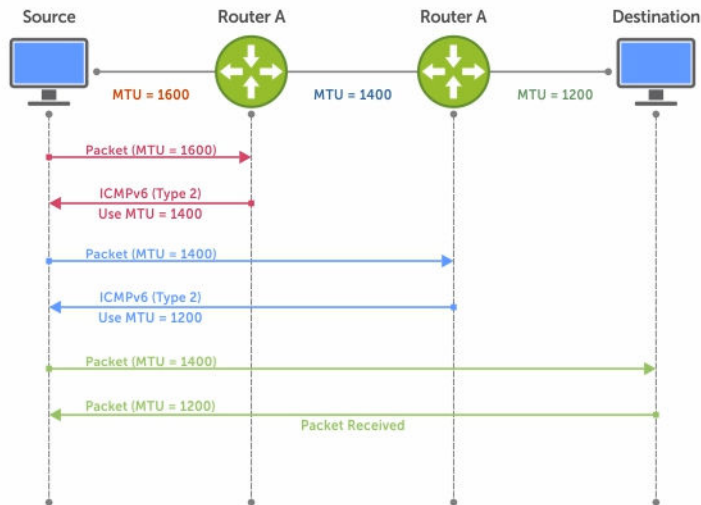




Figure 54. Path MTU Discovery Process

IPv6 Neighbor Discovery

IPv6 neighbor discovery protocol (NDP) is supported on the MXL switch platform.

NDP is a top-level protocol for neighbor discovery on an IPv6 network. In lieu of address resolution protocol (ARP), NDP uses "Neighbor Solicitation" and "Neighbor Advertisement" ICMPv6 messages for determining relationships between neighboring nodes. Using these messages, an IPv6 device learns the link-layer addresses for neighbors known to reside on attached links, quickly purging cached values that become invalid.

-  **NOTE:** If a neighboring node does not have an IPv6 address assigned, it must be manually pinged to allow the IPv6 device to determine the relationship of the neighboring node.
-  **NOTE:** To avoid problems with network discovery, Dell Networking recommends configuring the static route last or assigning an IPv6 address to the interface and assigning an address to the peer (the forwarding router's address) less than 10 seconds apart.

With ARP, each node broadcasts ARP requests on the entire link. This approach causes unnecessary processing by uninterested nodes. With NDP, each node sends a request only to the intended destination via a multicast address with the unicast address used as the last 24 bits. Other hosts on the link do not participate in the process, greatly increasing network bandwidth efficiency.




Figure 55. NDP Router Redirect

IPv6 Neighbor Discovery of MTU Packets

With the Dell Networking OS version 8.3.1.0, you can set the MTU advertised through the RA packets to incoming routers, without altering the actual MTU setting on the interface.

The `ipv6 nd mtu` command sets the value advertised to routers. It does not set the actual MTU rate. For example, if you set `ipv6 nd mtu` to 1280, the interface still passes 1500-byte packets, if that is what is set with the `mtu` command.

Configuring the IPv6 Recursive DNS Server

You can configure up to four Recursive DNS Server (RDNSS) addresses to be distributed via IPv6 router advertisements to an IPv6 device, using the `ipv6 nd dns-server ipv6-RDNSS-address {lifetime | infinite}` command in INTERFACE CONFIG mode.

The lifetime parameter configures the amount of time the IPv6 host can use the IPv6 RDNSS address for name resolution. The lifetime range is 0 to 4294967295 seconds. When the maximum lifetime value, 4294967295, or the `infinite` keyword is specified, the lifetime to use the RDNSS address does not expire. A value of 0 indicates to the host that the RDNSS address should not be used. You must specify a lifetime using the `lifetime` or `infinite` parameter.

The DNS server address does not allow the following:

- link local addresses
- loopback addresses
- prefix addresses
- multicast addresses

- invalid host addresses

If you specify this information in the IPv6 RDNSS configuration, a DNS error is displayed.

Example for Configuring an IPv6 Recursive DNS Server

The following example configures a RDNSS server with an IPv6 address of `1000:::1` and a lifetime of 1 second.

Debugging IPv6 RDNSS Information Sent to the Host

To verify that the IPv6 RDNSS information sent to the host is configured correctly, use the `debug ipv6 nd` command in EXEC Privilege mode.

Example of Debugging IPv6 RDNSS Information Sent to the Host

The following example debugs IPv6 RDNSS information sent to the host.

The last 3 lines indicate that the IPv6 RDNSS information was configured correctly.

```
dns-server=1000::0001, lifetime=1 sec
dns-server=3000::0001, lifetime=1 sec
dns-server=2000::0001, lifetime=0 sec
```

If the DNS server information is not displayed, verify that the IPv6 recursive DNS server configuration was configured on the correct interface.

Displaying IPv6 RDNSS Information

To display IPv6 interface information, including IPv6 RDNSS information, use the `show ipv6 interface` command in EXEC or EXEC Privilege mode.

Examples of Displaying IPv6 RDNSS Information

The following example displays IPv6 RDNSS information. The output in the last 3 lines indicates that the IPv6 RDNSS was correctly configured on interface `te 1/1`.

To display IPv6 RDNSS information, use the `show configuration` command in INTERFACE CONFIG mode.

The following example uses the `show configuration` command to display IPv6 RDNSS information.

IPv6 Multicast

IPv6 multicast is supported on the MXL switch platform.

The Dell Networking OS supports the following protocols to implement IPv6 multicast routing:

- **Multicast listener discovery protocol (MLD)** — MLD on a multicast router sends out periodic general MLD queries that the switch forwards through all ports in the VLAN. There are two versions of MLD: MLD version 1 is based on version 2 of the Internet group management protocol (IGMP) for IPv4; MLD version 2 is based on version 3 of the IGMP for IPv4. IPv6 multicast for the Dell Networking OS supports versions 1 and 2.
- **Protocol-independent multicast-sparse mode (PIM-SM)** — PIM-SM is a multicast protocol in which multicast receivers explicitly join to receive multicast traffic. The protocol uses a router as the root or

rendezvous point (RP) of the share tree distribution tree to distribute multicast traffic to a multicast group. Messages to join the multicast group (Join messages) are sent towards the RP and data is sent from senders to the RP so receivers can discover who are the senders and begin receiving traffic destined to the multicast group.

For more information, refer to the *Neighbor Discovery Protocol (NDP)*, *Multicast IPv6*, and *Protocol Independent Multicast (IPv6)* chapters in the *Dell Networking OS Command Line Interface Reference Guide*.

Secure Shell (SSH) Over an IPv6 Transport

IPv6 secure shell (SSH) is supported on the MXL switch platform.

The Dell Networking OS supports both inbound and outbound SSH sessions using IPv6 addressing. Inbound SSH supports accessing the system through the management interface as well as through a physical Layer 3 interface.

For SSH configuration details, refer to the *Security* chapter in the *Dell Networking OS Command Line Interface Reference Guide*.

Configuration Task List for IPv6

The following are configuration tasks for the IPv6 protocol.

- [Adjusting Your CAM-Profile](#)
- [Assigning an IPv6 Address to an Interface](#)
- [Assigning a Static IPv6 Route](#)
- [Configuring Telnet with IPv6](#)
- [SNMP over IPv6](#)
- [Showing IPv6 Information](#)
- [Clearing IPv6 Routes](#)

Adjusting Your CAM-Profile

The `cam-acl` command is supported on the MXL switch platform.

Although adjusting your CAM-profile is not a mandatory step, if you plan to implement IPv6 ACLs, adjust your CAM settings.

The CAM space is allotted in FP blocks. The total space allocated must equal 13 FP blocks. There are 16 FP blocks, but the System Flow requires three blocks that cannot be reallocated.

You must enter the `ipv6acl` allocation as a factor of 2 (2, 4, 6, 8, 10). All other profile allocations can use either even or odd-numbered ranges.

The default option sets the CAM Profile as follows:

- L3 ACL (ipv4acl): 6
- L2 ACL(l2acl): 5
- IPv6 L3 ACL (ipv6acl): 0

- L3 QoS (ipv4qos): 1
- L2 QoS (l2qos): 1

To have the changes take effect, save the new CAM settings to the startup-config (`write-mem` or `copy run start`) then reload the system for the new settings.

- Allocate space for IPV6 ACLs. Enter the CAM profile name then the allocated amount.

CONFIGURATION mode

```
cam-acl { ipv6acl }
```

When not selecting the default option, enter all of the profiles listed and a range for each.

The total space allocated must equal 13.

The `ipv6acl` range must be a factor of 2.

- Show the current CAM settings.

EXEC mode or EXEC Privilege mode

```
show cam-acl
```

- Provides information on FP groups allocated for the egress acl.

CONFIGURATION mode

```
show cam-acl-egress
```

Allocate at least one group for L2ACL and IPv4 ACL.

The total number of groups is 4.

Assigning an IPv6 Address to an Interface

IPv6 addresses are supported on the MXL switch platform.

Essentially, IPv6 is enabled in the Dell Networking OS simply by assigning IPv6 addresses to individual router interfaces. You can use IPv6 and IPv4 together on a system, but be sure to differentiate that usage carefully. To assign an IPv6 address to an interface, use the `ipv6 address` command.

You can configure up to two IPv6 addresses on management interfaces, allowing required default router support on the management port that is acting as host, per RFC 4861. Data ports support more than two IPv6 addresses.

When you configure IPv6 addresses on multiple interfaces (the `ipv6 address` command) and verify the configuration (the `show ipv6 interfaces` command), the same link local (fe80) address is displayed for each IPv6 interface.


- Enter the IPv6 Address for the device.

CONFIG-INTERFACE mode

```
ipv6 address ipv6 address/mask
```


– `ipv6 address`: x:x:x::x

– `mask`: The prefix length is from 0 to 128

 **NOTE:** IPv6 addresses are normally written as eight groups of four hexadecimal digits. Separate each group by a colon (:). Omitting zeros is accepted as described in [Addressing](#).

Assigning a Static IPv6 Route

IPv6 static routes are supported on the MXL switch platform.
To configure IPv6 static routes, use the `ipv6 route` command.

 **NOTE:** After you configure a static IPv6 route (the `ipv6 route` command) and configure the forwarding router's address (specified in the `ipv6 route` command) on a neighbor's interface, the IPv6 neighbor does not display in the `show ipv6 route` command output.

- Set up IPv6 static routes.
CONFIGURATION mode

```
ipv6 route prefix type {slot/port} forwarding router tag
- prefix: IPv6 route prefix
- type {slot/port}: interface type and slot/port
- forwarding router: forwarding router's address
- tag: route tag
```

Enter the keyword `interface` then the type of interface and slot/port information:

- For a 10/100/1000 Ethernet interface, enter the keyword `GigabitEthernet` then the slot/ port information.
- For a Gigabit Ethernet interface, enter the keyword `GigabitEthernet` then the slot/ port information.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
- For a Loopback interface, enter the keyword `loopback` then the loopback number.
- For a port-channel interface, enter the keywords `port-channel` then the port-channel number.
- For a VLAN interface, enter the keyword `vlan` then the VLAN ID.
- For a Null interface, enter the keyword `null` then the Null interface number.

Configuring Telnet with IPv6


IPv6 telnet is supported on the MXL switch platform.
The Telnet client and server in the Dell Networking OS supports IPv6 connections. You can establish a Telnet session directly to the router using an IPv6 Telnet client, or you can initiate an IPv6 Telnet connection from the router.

 **NOTE:** Telnet to link local addresses is supported on the MXL switch.

- Enter the IPv6 Address for the device.
EXEC mode or EXEC Privileged mode

```
telnet ipv6 address
- ipv6 address: x:x:x:x::x
```

- *mask*: prefix length is from 0 to 128.

 **NOTE:** IPv6 addresses are normally written as eight groups of four hexadecimal digits, where each group is separated by a colon (:). Omitting zeros is accepted as described in [Addressing](#).

SNMP over IPv6

The simple network management protocol (SNMP) is supported on the MXL switch platform.

You can configure SNMP over IPv6 transport so that an IPv6 host can perform SNMP queries and receive SNMP notifications from a device running IPv6. The Dell Networking OS SNMP-server commands for IPv6 have been extended to support IPv6. For more information regarding SNMP commands, refer to the *SNMP* and *SYSLOG* chapters in the *Dell Networking OS Command Line Interface Reference Guide*.

- `snmp-server host`
- `snmp-server user ipv6`
- `snmp-server community ipv6`
- `snmp-server community access-list-name ipv6`
- `snmp-server group ipv6`
- `snmp-server group access-list-name ipv6`

Showing IPv6 Information

All of the following `show` commands are supported on the MXL switch platform. View specific IPv6 configuration with the following commands.

- List the IPv6 show options.
EXEC mode or EXEC Privileged mode

```
show ipv6 ?
```

Example of `show ipv6` Command Options

```
Dell#show ipv6 ?
accounting  IPv6 accounting information
cam         IPv6 CAM Entries
fib         IPv6 FIB Entries
interface  IPv6 interface information
mbgproutes MBGP routing table
mld        MLD information
mroute     IPv6 multicast-routing table
neighbors  IPv6 neighbor information
ospf       OSPF information
pim        PIM V6 information
prefix-list List IPv6 prefix lists
route      IPv6 routing information
rpf        RPF table
Dell#
```

Showing an IPv6 Interface

To view the IPv6 configuration for a specific interface, use the following command.

- Show the currently running configuration for the specified interface.

EXEC mode

```
show ipv6 interface type {slot/port}
```

Enter the keyword `interface` then the type of interface and slot/port information:

- For all brief summary of IPv6 status and configuration, enter the keyword `brief`.
- For all IPv6 configured interfaces, enter the keyword `configured`.
- For a 10/100/1000 Ethernet interface, enter the keyword `GigabitEthernet` then the slot/ port information.
- For a Gigabit Ethernet interface, enter the keyword `GigabitEthernet` then the slot/ port information.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.
- For a loopback interface, enter the keyword `loopback` then the loopback number.
- For a port-channel interface, enter the keywords `port-channel` then the port-channel number.
- For a VLAN interface, enter the keyword `vlan` then the VLAN ID.

Showing IPv6 Routes

To view the global IPv6 routing information, use the following command.

- Show IPv6 routing information for the specified route type.

EXEC mode

```
show ipv6 route type
```

The following keywords are available:

- To display information about a network, enter `ipv6 address (X:X:X:X:X)`.
- To display information about a host, enter `hostname`.
- To display information about all IPv6 routes (including non-active routes), enter `all`.
- To display information about all connected IPv6 routes, enter `connected`.
- To display information about brief summary of all IPv6 routes, enter `summary`.
- To display information about Border Gateway Protocol (BGP) routes, enter `bgp`.
- To display information about ISO IS-IS routes, enter `isis`.
- To display information about Open Shortest Path First (OSPF) routes, enter `ospf`.
- To display information about Routing Information Protocol (RIP), enter `rip`.
- To display information about static IPv6 routes, enter `static`.
- To display information about an IPv6 Prefix lists, enter `list` and the prefix-list name.

Example of the show ipv6 route summary Command

Example of the show ipv6 route Command

Example of the show ipv6 route static Command

```
Dell#show ipv6 route summary
```

```
Route Source Active Routes Non-active Routes
connected    5             0
static       0             0
Total        5             0
```

```
Dell#show ipv6 route
```

```
Codes: C - connected, L - local, S - static, R - RIP,
        B - BGP, IN - internal BGP, EX - external BGP, LO - Locally Originated,
        O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
        N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
        E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,
        L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,
Gateway of last resort is not set
```

```
Destination Dist/Metric, Gateway, Last Change
-----
C 2001::/64 [0/0]
   Direct, Gi 1/1, 00:28:49
C 2002::/120 [0/0]
   Direct, Gi 1/1, 00:28:49
C 2003::/120 [0/0]
   Direct, Gi 1/1, 00:28:49
```

```
Dell#show ipv6 route static
```

```
Destination Dist/Metric, Gateway, Last Change
-----
S          8888:9999:5555:6666:1111:2222::/96 [1/0]
   via          2222:2222:3333:3333::1, Gi 9/1, 00:03:16
S          9999:9999:9999:9999::/64 [1/0]
   via 8888:9999:5555:6666:1111:2222:3333:4444, 00:03:16
```

Showing the Running-Configuration for an Interface

To view the configuration for any interface, use the following command.

- Show the currently running configuration for the specified interface.

EXEC mode

```
show running-config interface type {slot/port}
```

Enter the keyword `interface` then the type of interface and slot/port information:

- For a 10/100/1000 Ethernet interface, enter the keyword `GigabitEthernet` then the slot/ port information.
- For a Gigabit Ethernet interface, enter the keyword `GigabitEthernet` then the slot/ port information.
- For the Management interface on the RPM, enter the keyword `ManagementEthernet` then the slot/port information.

- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

Example of the `show running-config interface` Command

```
Dell#show run int gi 2/2
!
interface GigabitEthernet 2/2
  no ip address
  ipv6 address 3:4:5:6::8/24
  shutdown
Dell#
```


Clearing IPv6 Routes

To clear routes from the IPv6 routing table, use the following command.

- Clear (refresh) all or a specific route from the IPv6 routing table.
EXEC mode

```
clear ipv6 route {* | ipv6 address prefix-length}
```

- *: all routes.
- *ipv6 address*: the format is `x:x:x::x`.
- *mask*: the prefix length is from 0 to 128.

 **NOTE:** IPv6 addresses are normally written as eight groups of four hexadecimal digits, where each group is separated by a colon (:). Omitting zeros is accepted as described in [Addressing](#).

iSCSI Optimization

The MXL switch enables internet small computer system interface (iSCSI) optimization with default iSCSI parameter settings and is auto-provisioned to support the following features.

- [Detection and Auto-Configuration for Dell EqualLogic Arrays](#)
- [Configuring Detection and Ports for Dell Compellent Arrays](#)

To display information on iSCSI configuration and sessions, use the `show` commands.

iSCSI optimization enables quality-of-service (QoS) treatment for iSCSI traffic.

iSCSI Optimization Overview

iSCSI is a TCP/IP-based protocol for establishing and managing connections between IP-based storage devices and initiators in a storage area network (SAN).

iSCSI optimization enables the network switch to auto-detect Dell's iSCSI storage arrays and triggers a self-configuration of several key network configurations that enables optimization of the network for better storage traffic throughput.

iSCSI optimization also provides a means of monitoring iSCSI sessions and applying quality of service (QoS) policies on iSCSI traffic. When enabled, iSCSI optimization allows a switch to monitor (snoop) the establishment and termination of iSCSI connections. The switch uses the snooped information to detect iSCSI sessions and connections established through the switch.

iSCSI optimization allows you to reduce deployment time and management complexity in data centers. In a data center network, Dell EqualLogic and Compellent iSCSI storage arrays are connected to a converged Ethernet network using the data center bridging exchange protocol (DCBx) through stacked and/or non-stacked Ethernet switches.

iSCSI session monitoring over virtual link trunking (VLT) synchronizes the iSCSI session information between the VLT peers, allowing session information to be available in both the VLT peers.

iSCSI optimization functions as follows:

- Auto-detection of EqualLogic storage arrays — the switch detects any active EqualLogic array directly attached to its ports.
- Manual configuration to detect Compellent storage arrays where auto-detection is not supported.
- Automatic configuration of switch ports after detection of storage arrays.
- If you configure flow-control, iSCSI uses the current configuration. If you do not configure flow-control, iSCSI auto-configures flow control.

- iSCSI monitoring sessions — the switch monitors and tracks active iSCSI sessions in connections on the switch, including port information and iSCSI session information.
- iSCSI QoS — A user-configured iSCSI class of service (CoS) profile is applied to all iSCSI traffic. Classifier rules are used to direct the iSCSI data traffic to queues that can be given preferential QoS treatment over other data passing through the switch. Preferential treatment helps to avoid session interruptions during times of congestion that would otherwise cause dropped iSCSI packets.
- iSCSI DCBx TLVs are supported.

The following illustration shows iSCSI optimization between servers in an M1000e enclosure and a storage array in which a stack and MXL connects installed servers (iSCSI initiators) to a storage array (iSCSI targets) in a SAN network. iSCSI optimization running on the MXL is configured to use dot1p priority-queue assignments to ensure that iSCSI traffic in these sessions receives priority treatment when forwarded on MXL hardware.

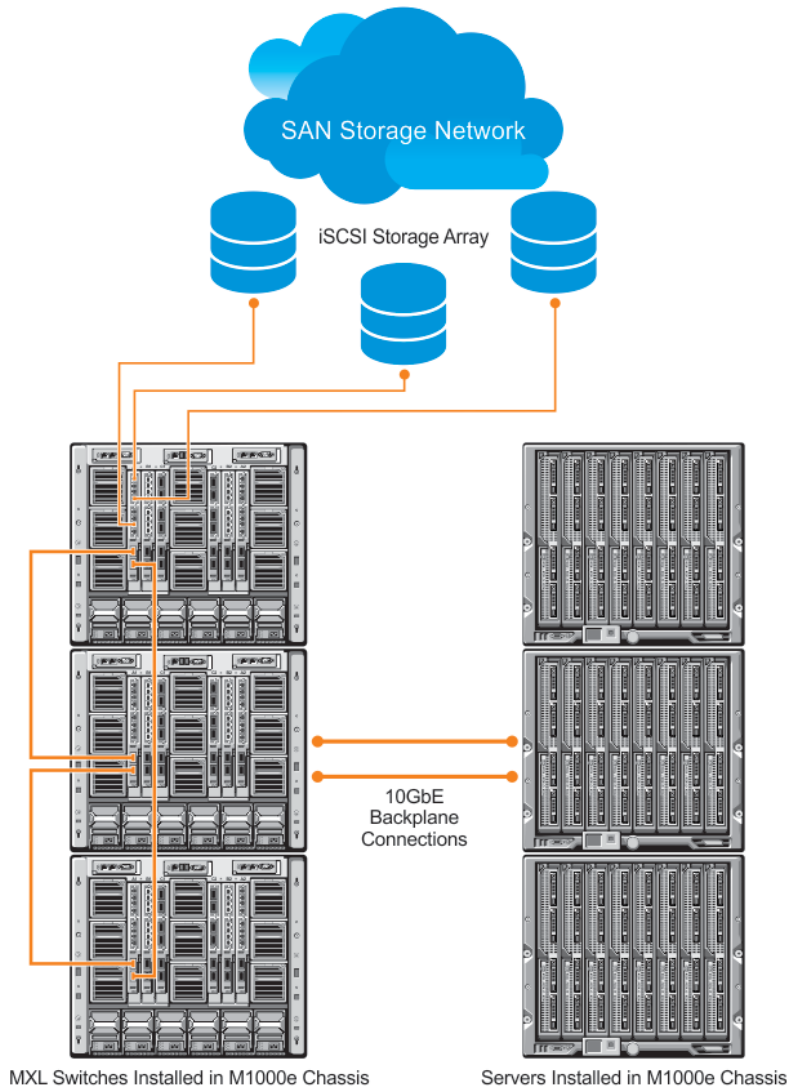


Figure 56. iSCSI Optimization Example

Monitoring iSCSI Traffic Flows

The switch snoops iSCSI session-establishment and termination packets by installing classifier rules that trap iSCSI protocol packets to the CPU for examination.

Devices that initiate iSCSI sessions usually use well-known TCP ports 3260 or 860 to contact targets. When you enable iSCSI optimization, by default the switch identifies IP packets to or from these ports as iSCSI traffic.

You can configure the switch to monitor traffic for additional port numbers or a combination of port number and target IP address, and you can remove the well-known port numbers from monitoring.

Information Monitored in iSCSI Traffic Flows

iSCSI optimization examines the following data in packets and uses the data to track the session and create the classifier entries that enable QoS treatment.

- Initiator's IP Address
- Target's IP Address
- Initiator defined session identifier (ISID)
- Initiator's iSCSI qualified name (IQN)
- Target's IQN
- Initiator's TCP Port
- Target's TCP Port

If no iSCSI traffic is detected for a session during a user-configurable aging period, the session data is cleared.

Detection and Auto-Configuration for Dell EqualLogic Arrays

The iSCSI optimization feature includes auto-provisioning support with the ability to detect directly connected Dell EqualLogic storage arrays and automatically reconfigure the switch to enhance storage traffic flows.

The MXL uses the link layer discovery protocol (LLDP) to discover Dell EqualLogic devices on the network. LLDP is enabled by default. For more information about LLDP, refer to [Link Layer Discovery Protocol \(LLDP\)](#).

The following message displays the first time a Dell EqualLogic array is detected and describes the configuration changes that are automatically performed:

```
%STKUNIT0-M:CP %IFMGR-5-IFM_ISCSI_AUTO_CONFIG: This switch is being configured
for optimal
conditions to support iSCSI traffic which will cause some automatic
configuration to occur
including jumbo frames and flow-control on all ports; no storm control and
spanning-tree port
fast to be enabled on the port of detection.
```

The following syslog message is generated the first time an EqualLogic array is detected:

```
%STKUNIT0-M:CP %LLDP-5-LLDP_EQL_DETECTED: EqualLogic Storage Array detected on
interface Te 1/43
```

- At the first detection of an EqualLogic array, an MTU of 12000 is enabled on all ports and port-channels (if it has not already been enabled).
- Spanning-tree portfast is enabled on the interface LLDP identifies.
- Unicast storm control is disabled on the interface LLDP identifies.

Configuring Detection and Ports for Dell Compellent Arrays

For the best iSCSI traffic conditions, the MXL switch auto-configures a port connected to a Dell Compellent storage array, when configured as compellent connected port through CLI.

The following message displays the first time a Dell Compellent storage array is detected and describes the configuration changes that are automatically performed:

```
%STKUNIT0-M:CP %IFMGR-5-IFM_ISCSI_AUTO_CONFIG: This switch is being configured
for optimal
conditions to support iSCSI traffic which will cause some automatic
configuration to occur
including jumbo frames and flow-control on all ports; no storm control and
spanning-tree port
fast to be enabled on the port of detection.
```

The MXL switch auto-configures as follows:

- At the first detection, an MTU of 12000 is enabled on all ports and port-channels (if it is not already enabled).
- Spanning-tree portfast is enabled on the interface identified by CLI, if the port is in L2 mode.
- Unicast storm control is disabled on the interface identified by CLI.

iSCSI Optimization: Operation

iSCSI optimization requires LLDP to be enabled. LLDP is enabled by default on MXL switch.

When the MXL auto-configures with iSCSI enabled, the following actions occurs:

- Link-level flow control is enabled on PFC disabled interfaces.
- iSCSI session snooping is enabled.
- iSCSI LLDP monitoring starts to automatically detect EqualLogic arrays.

The following message displays when you enable iSCSI on a switch and describes the configuration changes that are automatically performed:

```
%STKUNIT0-M:CP %IFMGR-5-IFM_ISCSI_ENABLE: iSCSI has been enabled causing flow
control to be enabled on all interfaces. EQL detection and enabling iscsi
profile-compellent on an interface may cause some automatic configurations to
occur like jumbo frames on all ports and no storm control and spanning tree
port-fast on the port of detection.
```

Default iSCSI Optimization Values

The following table lists the default values for the iSCSI optimization feature.

Table 26. iSCSI Optimization Defaults

Parameter	Default Value
iSCSI Optimization global setting	Enabled
iSCSI CoS mode (802.1p priority queue mapping)	Enabled: dot1p priority 4 without the <code>remark</code> setting
iSCSI CoS Treatment	iSCSI packets are queued based on dot1p instead of DSCP values.
VLAN priority tag	iSCSI flows are assigned by default to dot1p priority 4 without the <code>remark</code> setting.
DSCP	None: user-configurable.
Remark	Not configured.
iSCSI session aging time	10 minutes
iSCSI optimization target ports	iSCSI well-known ports 3260 and 860 are configured as default (with no IP address or name) but can be removed as any other configured target.
iSCSI session monitoring	Enabled. The CAM allocation for iSCSI is set to two.

Displaying iSCSI Optimization Information

To display information on iSCSI optimization, use the following `show` commands.

- Display the currently configured iSCSI settings.
`show iscsi`
- Display information on active iSCSI sessions on the switch.
`show iscsi sessions`
- Display detailed information on active iSCSI sessions on the switch. To display detailed information on specified iSCSI session, enter the session's iSCSI ID.
`show iscsi sessions detailed [session isid]`
- Display all globally configured non-default iSCSI settings in the current session.
`show run iscsi`

Example of the `show iscsi` Command

Example of the `show iscsi session` Command

Example of the `show iscsi session detailed` Command

```
Dell#show iscsi
iSCSI is enabled
iSCSI session monitoring is disabled
```

```
iSCSI COS : dotlp is 4 no-remark
Session aging time: 10
Maximum number of connections is 256
```

```
-----
iSCSI Targets and TCP Ports:
```

```
-----
TCP Port Target IP Address
3260
860
```

```
VLT PEER1
```

```
Dell#show iscsi session
Session 0:
```

```
-----
Target: iqn.2001-05.com.equallogic:0-8a0906-0e70c2002-10a0018426a48c94-iom010
Initiator: iqn.1991-05.com.microsoft:win-x918v27yajg
ISID: 400001370000
VLT PEER2
```

```
Session 1:
```

```
-----
Target: iqn.2001-05.com.equallogic:0-8a0906-0f60c2002-0360018428d48c94-iom011
iqn.1991-05.com.microsoft:win-x918v27yajg
ISID: 400001370000
```

```
Dell# show iscsi sessions detailed
Session 0 :
```

```
-----
Target:iqn.2010-11.com.ixia:ixload:iscsi-TG1
Initiator:iqn.2010-11.com.ixia.ixload:initiator-iscsi-2c
Up Time:00:00:01:28(DD:HH:MM:SS)
Time for aging out:00:00:09:34(DD:HH:MM:SS)
ISID:806978696102
Initiator Initiator Target Target Connection
IP Address TCP Port IP Address TCPPort ID
10.10.0.44 33345 10.10.0.101 3260 0
Session 1 :
```

```
-----
Target:iqn.2010-11.com.ixia:ixload:iscsi-TG1
Initiator:iqn.2010-11.com.ixia.ixload:initiator-iscsi-35
Up Time:00:00:01:22(DD:HH:MM:SS)
Time for aging out:00:00:09:31(DD:HH:MM:SS)
ISID:806978696102
Initiator Initiator Target Target Connection
IP Address TCP Port IP Address TCPPort ID
10.10.0.53 33432 10.10.0.101 3260 0
```


Intermediate System to Intermediate System

Intermediate system to intermediate system (IS-IS) is supported on the MXL switch platform.

- The IS-IS protocol is an interior gateway protocol (IGP) that uses a shortest-path-first algorithm. Dell Networking supports both IPv4 and IPv6 versions of IS-IS.
- The IS-IS protocol standards are listed in the [Standards Compliance](#) chapter.

IS-IS Protocol Overview

The IS-IS protocol, developed by the International Organization for Standardization (ISO), is an interior gateway protocol (IGP) that uses a shortest-path-first algorithm.

 **NOTE:** This protocol supports routers passing both IP and OSI traffic, though the Dell Networking implementation supports only IP traffic.

IS-IS uses the following management information base (MIB): `draft-ietf-isis-wg-mib-16` and `f10-isis`.

IS-IS is organized hierarchically into routing domains and each router or system resides in at least one area. In IS-IS, routers are designated as Level 1, Level 2 or Level 1-2 systems. Level 1 routers only route traffic within an area, while Level 2 routers route traffic between areas. At its most basic, Level 1 systems route traffic within the area and any traffic destined for outside the area is sent to a Level 1-2 system. Level 2 systems manage destination paths for external routers. Only Level 2 routers can exchange data packets or routing information directly with external routers located outside of the routing domains. Level 1-2 systems manage both inter-area and intra-area traffic by maintaining two separate link databases; one for Level 1 routes and one for Level 2 routes. A Level 1-2 router does not advertise Level 2 routes to a Level 1 router.

To establish adjacencies, each IS-IS router sends different protocol data units (PDU). For IP traffic, the IP addressing information is included in the IS-IS hello PDUs and the link state PDUs (LSPs).

This brief overview is not intended to provide a complete understanding of IS-IS; for that, consult the documents listed in [Multi-Topology IS-IS](#).

IS-IS Addressing

IS-IS PDUs require ISO-style addressing called network entity title (NET).

For those familiar with name-to-network service mapping point (NSAP) addresses, the composition of the NET is identical to an NSAP address, except the last byte is always 0. The NET is composed of the IS-

IS area address, system ID, and N-selector. The last byte is the N-selector. All routers within an area have the same area portion. Level 1 routers route based on the system address portion of the address, while the Level 2 routers route based on the area address.

The NET length is variable, with a maximum of 20 bytes and a minimum of 8 bytes. It is composed of the following:

- **area address** — within your routing domain or area, each area must have a unique area value. The first byte is called the authority and format indicator (AFI).
- **system address** — the router's MAC address.
- **N-selector** — this is always 0.

The following illustration is an example of the ISO-style address to show the address format IS-IS uses. In this example, the first five bytes (47.0005.0001) are the area address. The system portion is 000c.000a.4321 and the last byte is always 0.



Figure 57. ISO Address Format

Multi-Topology IS-IS

Multi-topology IS-IS (MT IS-IS) allows you to create multiple IS-IS topologies on a single router with separate databases.

Use this feature to place a virtual physical topology into logical routing domains, which can each support different routing and security policies.

All routers on a LAN or point-to-point must have at least one common supported topology when operating in Multi-Topology IS-IS mode. If IPv4 is the common supported topology between those two routers, adjacency can be formed. All topologies must share the same set of L1-L2 boundaries.

You must implement a wide metric-style globally on the autonomous system (AS) to run multi-topology IS-IS for IPv6 because the Type, Length, Value (TLVs) used to advertise IPv6 information in link-state packets (LSPs) are defined to use only extended metrics.

The multi-topology ID is shown in the first octet of the IS-IS packet. Certain MT topologies are assigned to serve predetermined purposes:

- MT ID #0: Equivalent to the "standard" topology.
- MT ID #1: Reserved for IPv4 in-band management purposes.
- MT ID #2: Reserved for IPv6 routing topology.
- MT ID #3: Reserved for IPv4 multicast routing topology.
- MT ID #4: Reserved for IPv6 multicast routing topology.

- MT ID #5: Reserved for IPv6 in-band management purposes.

Transition Mode

All routers in the area or domain must use the same type of IPv6 support, either single-topology or multi-topology. A router operating in multi-topology mode does not recognize the ability of the single-topology mode router to support IPv6 traffic, which leads to holes in the IPv6 topology.

While in Transition mode, both types of TLVs (single-topology and multi-topology) are sent in LSPs for all configured IPv6 addresses, but the router continues to operate in single-topology mode (that is, the topological restrictions of the single-topology mode remain in effect). Transition mode stops after all routers in the area or domain have been upgraded to support multi-topology IPv6. After all routers in the area or domain are operating in multi-topology IPv6 mode, the topological restrictions of single-topology mode are no longer in effect.

Interface Support

MT IS-IS is supported on physical Ethernet interfaces, physical synchronous optical network technologies (SONET) interfaces, port-channel interfaces (static and dynamic using LACP), and virtual local area network (VLAN) interfaces.

Adjacencies

Adjacencies on point-to-point interfaces are formed as usual, where IS-IS routers do not implement MT extensions.

If a local router does not participate in certain MTs, it does not advertise those MT IDs in its IS-IS hellos (IIHs) and so does not include that neighbor within its LSPs. If an MT ID is not detected in the remote side's IIHs, the local router does not include that neighbor within its LSPs. The local router does not form an adjacency if both routers do not have at least one common MT over the interface.

Graceful Restart

Graceful Restart is supported on MXL platforms for both Helper and Restart modes.

Graceful restart is a protocol-based mechanism that preserves the forwarding table of the restarting router and its neighbors for a specified period to minimize the loss of packets. A graceful-restart router does not immediately assume that a neighbor is permanently down and so does not trigger a topology change.

Normally, when an IS-IS router is restarted, temporary disruption of routing occurs due to events in both the restarting router and the neighbors of the restarting router. When a router goes down without a graceful restart, there is a potential to lose access to parts of the network due to the necessity of network topology changes.

IS-IS graceful restart recognizes the fact that in a modern router, the control plane and data plane are functionally separate. Restarting the control plane functionality (such as the failover of the active route processor module (RPM) to the backup in a redundant configuration) should not necessarily interrupt data packet forwarding. This behavior is supported because the forwarding tables previously computed

by an active RPM have been downloaded into the forwarding information base (FIB) on the line cards (the data plane) and are still resident. For packets that have existing FIB/content addressable memory (CAM) entries, forwarding between ingress and egress ports can continue uninterrupted while the control plane IS-IS process comes back to full functionality and rebuilds its routing tables.

A new TLV (the Restart TLV) is introduced in the IIH PDUs, indicating that the router supports graceful restart.

Timers

Three timers are used to support IS-IS graceful restart functionality. After you enable graceful restart, these timers manage the graceful restart process.

There are three times, T1, T2, and T3.

- The T1 timer specifies the wait time before unacknowledged restart requests are generated. This is the interval before the system sends a Restart Request (an IIH with the RR bit set in Restart TLV) until the complete sequence number PDU (CSNP) is received from the helping router. You can set the duration to a specific amount of time (seconds) or a number of attempts.
- The T2 timer is the maximum time that the system waits for LSP database synchronization. This timer applies to the database type (level-1, level-2, or both).
- The T3 timer sets the overall wait time after which the router determines that it has failed to achieve database synchronization (by setting the overload bit in its own LSP). You can base this timer on adjacency settings with the value derived from adjacent routers that are engaged in graceful restart recovery (the minimum of all the Remaining Time values advertised by the neighbors) or by setting a specific amount of time manually.

Implementation Information

IS-IS implementation supports one instance of IS-IS and six areas.

You can configure the system as a Level 1 router, a Level 2 router, or a Level 1-2 router. For IPv6, the IPv4 implementation has been expanded to include two new type, length, values (TLVs) in the PDU that carry information required for IPv6 routing. The new TLVs are *IPv6 Reachability* and *IPv6 Interface Address*. Also, a new IPv6 protocol identifier has also been included in the supported TLVs. The new TLVs use the extended metrics and up/down bit semantics.

Multi-topology IS-IS adds TLVs:

- **MT TLV** — contains one or more Multi-Topology IDs in which the router participates. This TLV is included in IIH and the first fragment of an LSP.
- **MT Intermediate Systems TLV** — appears for every topology a node supports. An MT ID is added to the extended IS reachability TLV type 22.
- **MT Reachable IPv4 Prefixes TLV** — appears for each IPv4 an IS announces for a given MT ID. Its structure is aligned with the extended IS Reachability TLV Type 236 and it adds an MT ID.
- **MT Reachable IPv6 Prefixes TLV** — appears for each IPv6 an IS announces for a given MT ID. Its structure is aligned with the extended IS Reachability TLV Type 236 and add an MT ID.

By default, the system supports dynamic host name exchange to assist with troubleshooting and configuration. By assigning a name to an IS-IS NET address, you can track IS-IS information on that address easier. The system does not support ISO CLNS routing; however, the ISO NET format is supported for addressing.

To support IPv6, the Dell Networking implementation of IS-IS performs the following tasks:

- Advertises IPv6 information in the PDUs.
- Processes IPv6 information received in the PDUs.
- Computes routes to IPv6 destinations.
- Downloads IPv6 routes to the RTM for installing in the FIB.
- Accepts external IPv6 information and advertises this information in the PDUs.


The following table lists the default IS-IS values.

Table 27. IS-IS Default Values

IS-IS Parameter	Default Value
Complete sequence number PDU (CSNP) interval	10 seconds
IS-to-IS hello PDU interval	10 seconds
IS-IS interface metric	10
Metric style	Narrow
Designated Router priority	64
Circuit Type	Level 1 and Level 2
IS Type	Level 1 and Level 2
Equal Cost Multi Paths	16

Configuration Information

To use IS-IS, you must configure and enable IS-IS in two or three modes: CONFIGURATION ROUTER ISIS, CONFIGURATION INTERFACE, and (when configuring for IPv6) ADDRESS-FAMILY mode. Commands in ROUTER ISIS mode configure IS-IS globally, while commands executed in INTERFACE mode enable and configure IS-IS features on that interface only. Commands in the ADDRESS-FAMILY mode are specific to IPv6.

 **NOTE:** When using the IS-IS routing protocol to exchange IPv6 routing information and to determine destination reachability, you can route IPv6 along with IPv4 while using a single intra-domain routing protocol. The configuration commands allow you to enable and disable IPv6 routing and to configure or remove IPv6 prefixes on links.

Except where identified, the commands described in this chapter apply to both IPv4 and IPv6 versions of IS-IS.

Configuration Tasks for IS-IS

The following describes the configuration tasks for IS-IS.

- [Enabling IS-IS](#)
- [Configuring Multi-Topology IS-IS \(MT IS-IS\)](#)
- [Configuring IS-IS Graceful Restart](#)
- [Changing LSP Attributes](#)


- [Configuring the IS-IS Metric Style](#)
- [Configuring the IS-IS Cost](#)
- [Changing the IS-Type](#)
- [Controlling Routing Updates](#)
- [Configuring Authentication Passwords](#)
- [Setting the Overload Bit](#)
- [Debugging IS-IS](#)

Enabling IS-IS

By default, IS-IS is not enabled.

The system supports one instance of IS-IS. To enable IS-IS globally, create an IS-IS routing process and assign a NET address. To exchange protocol information with neighbors, enable IS-IS on an interface, instead of on a network as with other routing protocols.

In IS-IS, neighbors form adjacencies only when they are same IS type. For example, a Level 1 router never forms an adjacency with a Level 2 router. A Level 1-2 router forms Level 1 adjacencies with a neighboring Level 1 router and forms Level 2 adjacencies with a neighboring Level 2 router.

 **NOTE:** Even though you enable IS-IS globally, enable the IS-IS process on an interface for the IS-IS process to exchange protocol information and form adjacencies.

To configure IS-IS globally, use the following commands.

1. Create an IS-IS routing process.

CONFIGURATION mode

```
router isis [tag]
```

tag: (optional) identifies the name of the IS-IS process.

2. Configure an IS-IS network entity title (NET) for a routing process.

ROUTER ISIS mode

```
net network-entity-title
```

Specify the area address and system ID for an IS-IS routing process. The last byte must be 00.

For more information about configuring a NET, refer to [IS-IS Addressing](#).

3. Enter the interface configuration mode.

CONFIGURATION mode

```
interface interface
```

Enter the keyword *interface* then the type of interface and slot/port information:

- For the Loopback interface on the RPM, enter the keyword `loopback` then a number from 0 to 16383.
- For a port channel, enter the keywords `port-channel` then a number from 1 to 255.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `FortyGigabitEthernet` then the slot/port information.

- For a VLAN, enter the keyword `vlan` then a number from 1 to 4094.
4. Enter an IPv4 Address.

INTERFACE mode

```
ip address ip-address mask
```

Assign an IP address and mask to the interface.

The IP address must be on the same subnet as other IS-IS neighbors, but the IP address does not need to relate to the NET address.

5. Enter an IPv6 Address.

INTERFACE mode

```
ipv6 address ipv6-address mask
```

- `ipv6 address: x:x:x:x::x`
- `mask`: The prefix length is from 0 to 128.

The IPv6 address must be on the same subnet as other IS-IS neighbors, but the IP address does not need to relate to the NET address.

6. Enable IS-IS on the IPv4 interface.

ROUTER ISIS mode

```
ip router isis [tag]
```

If you configure a tag variable, it must be the same as the `tag` variable assigned in step 1.

7. Enable IS-IS on the IPv6 interface.

ROUTER ISIS mode

```
ipv6 router isis [tag]
```

If you configure a tag variable, it must be the same as the `tag` variable assigned in step 1.

Example of Viewing IS-IS Configuration (EXEC Privilege Mode)

Example of the `show isis traffic` Command

The default IS type is **level-1-2**. To change the IS type to Level 1 only or Level 2 only, use the `is-type` command in ROUTER ISIS mode.

To view the IS-IS configuration, enter the `show isis protocol` command in EXEC Privilege mode or the `show config` command in ROUTER ISIS mode.

```
Dell#show isis protocol
IS-IS Router: <Null Tag>
System Id: EEEE.EEEE.EEEE IS-Type: level-1-2
Manual area address(es):
 47.0004.004d.0001
Routing for area address(es):
 21.2223.2425.2627.2829.3031.3233
 47.0004.004d.0001
Interfaces supported by IS-IS:
 Vlan 2
 GigabitEthernet 4/22
 Loopback 0
```

```
Redistributing:
Distance: 115
Generate narrow metrics: level-1-2
Accept narrow metrics:   level-1-2
Generate wide metrics:   none
Accept wide metrics:     none
Dell#
```

To view IS-IS protocol statistics, use the `show isis traffic` command in EXEC Privilege mode.

```
Dell#show isis traffic
IS-IS: Level-1 Hellos (sent/rcvd) : 4272/1538
IS-IS: Level-2 Hellos (sent/rcvd) : 4272/1538
IS-IS: PTP Hellos (sent/rcvd) : 0/0
IS-IS: Level-1 LSPs sourced (new/refresh) : 0/0
IS-IS: Level-2 LSPs sourced (new/refresh) : 0/0
IS-IS: Level-1 LSPs flooded (sent/rcvd) : 32/19
IS-IS: Level-2 LSPs flooded (sent/rcvd) : 32/17
IS-IS: Level-1 LSPs CSNPs (sent/rcvd) : 1538/0
IS-IS: Level-2 LSPs CSNPs (sent/rcvd) : 1534/0
IS-IS: Level-1 LSPs PSNPs (sent/rcvd) : 0/0
IS-IS: Level-2 LSPs PSNPs (sent/rcvd) : 0/0
IS-IS: Level-1 DR Elections : 2
IS-IS: Level-2 DR Elections : 2
IS-IS: Level-1 SPF Calculations : 29
IS-IS: Level-2 SPF Calculations : 29
IS-IS: LSP checksum errors received : 0
IS-IS: LSP authentication failures : 0
Dell#
```

You can assign more NET addresses, but the System ID portion of the NET address must remain the same. The Dell Networking OS supports up to six area addresses.

Some address considerations are:

- In order to be neighbors, configure Level 1 routers with at least one common area address.
- A Level 2 router becomes a neighbor with another Level 2 router regardless of the area address configured. However, if the area addresses are different, the link between the Level 2 routers is only at Level 2.

Configuring Multi-Topology IS-IS (MT IS-IS)

To configure multi-topology IS-IS (MT IS-IS), use the following commands.

1. Enable multi-topology IS-IS for IPv6.

```
ROUTER ISIS AF IPV6 mode
```

```
multi-topology [transition]
```

Enter the keyword *transition* to allow an IS-IS IPv6 user to continue to use single-topology mode while upgrading to multi-topology mode. After every router has been configured with the *transition* keyword, and all the routers are in MT IS-IS IPv6 mode, you can remove the *transition* keyword on each router.



NOTE: When you do not enable transition mode, you do not have IPv6 connectivity between routers operating in single-topology mode and routers operating in multi-topology mode.

2. Exclude this router from other router's SPF calculations.

ROUTER ISIS AF IPV6 mode

```
set-overload-bit
```

3. Set the minimum interval between SPF calculations.

ROUTER ISIS AF IPV6 mode

```
spf-interval [level-1 | level-2 | interval] [initial_wait_interval  
[second_wait_interval]]
```

Use this command for IPv6 route computation only when you enable multi-topology. If using Single-Topology mode, to apply to both IPv4 and IPv6 route computations, use the `spf-interval` command in CONFIG ROUTER ISIS mode.

4. Implement a *wide metric-style* globally.

ROUTER ISIS AF IPV6 mode

```
isis ipv6 metric metric-value [level-1 | level-2 | level-1-2]
```

To configure wide or wide transition metric style, the cost can be between 0 and 16,777,215.

Configuring IS-IS Graceful Restart

To enable IS-IS graceful restart globally, use the following commands. Additionally, you can implement optional commands to enable the graceful restart settings.

- Enable graceful restart on ISIS processes.

ROUTER-ISIS mode

```
graceful-restart ietf
```

- Configure the time during which the graceful restart attempt is prevented.

ROUTER-ISIS mode

```
graceful-restart interval minutes
```

The range is from 1 to 120 minutes.

The default is **5 minutes**.

- Enable the graceful restart maximum wait time before a restarting peer comes up.

ROUTER-ISIS mode

```
graceful-restart restart-wait seconds
```

When implementing this command, be sure to set the T3 timer to adjacency on the restarting router.

The range is from 1 to 120 minutes.

The default is **30 seconds**.

- Configure the time that the graceful restart timer T1 defines for a restarting router to use for each interface, as an interval before regenerating Restart Request (an IIH with RR bit set in Restart TLV) after waiting for an acknowledgement.

ROUTER-ISIS mode

```
graceful-restart t1 {interval seconds | retry-times value}
```

- `interval`: wait time (the range is from 5 to 120. The default is **5**.)
- `retry-times`: number of times an unacknowledged restart request is sent before the restarting router gives up the graceful restart engagement with the neighbor. (The range is from 1 to 10 attempts. The default is **1**.)
- Configure the time for the graceful restart timer T2 that a restarting router uses as the wait time for each database to synchronize.

ROUTER-ISIS mode

```
graceful-restart t2 {level-1 | level-2} seconds
```

- `level-1, level-2`: identifies the database instance type to which the wait interval applies.

The range is from 5 to 20 seconds.

The default is **30 seconds**.

- Configure graceful restart timer T3 to set the time used by the restarting router as an overall maximum time to wait for database synchronization to complete.

ROUTER-ISIS mode

```
graceful-restart t3 {adjacency | manual seconds}
```


- `adjacency`: the restarting router receives the remaining time value from its peer and adjusts its T3 value so if user has configured this option.
- `manual`: allows you to specify a fixed value that the restarting router should use.

The range is from 50 to 120 seconds.

The default is **30 seconds**.

Example of the `show isis graceful-restart detail` Command

Example of the `show isis interface` Command

 **NOTE:** If this timer expires before the synchronization has completed, the restarting router sends the `overload` bit in the LSP. The `overload` bit is an indication to the receiving router that database synchronization did not complete at the restarting router.

To view all graceful restart-related configurations, use the `show isis graceful-restart detail` command in EXEC Privilege mode.

```
Dell#show isis graceful-restart detail
Configured Timer Value
=====
Graceful Restart      : Enabled
Interval/Blackout time : 1 min
T3 Timer              : Manual
T3 Timeout Value      : 30
T2 Timeout Value      : 30 (level-1), 30 (level-2)
T1 Timeout Value      : 5, retry count: 1
Adjacency wait time   : 30

Operational Timer Value
=====
Current Mode/State    : Normal/RUNNING
T3 Time left          : 0
T2 Time left          : 0 (level-1), 0 (level-2)
Restart ACK rcv count : 0 (level-1), 0 (level-2)
Restart Req rcv count : 0 (level-1), 0 (level-2)
```

```
Suppress Adj rcv count : 0 (level-1), 0 (level-2)
Restart CSNP rcv count : 0 (level-1), 0 (level-2)
Database Sync count    : 0 (level-1), 0 (level-2)
```

```
Circuit GigabitEthernet 2/10:
  Mode: Normal L1-State:NORMAL, L2-State: NORMAL

  L1: Send/Receive: RR:0/0, RA: 0/0, SA:0/0
      T1 time left: 0, retry count left:0

  L2: Send/Receive: RR:0/0, RA: 0/0, SA:0/0
      T1 time left: 0, retry count left:0
Dell#
```

To view all interfaces configured with IS-IS routing along with the defaults, use the `show isis interface` command in EXEC Privilege mode.

```
Dell#show isis interface G1/34
GigabitEthernet 2/10 is up, line protocol is up
  MTU 1497, Encapsulation SAP
  Routing Protocol: IS-IS
    Circuit Type: Level-1-2
    Interface Index 0x62cc03a, Local circuit ID 1
    Level-1 Metric: 10, Priority: 64, Circuit ID: 0000.0000.000B.01
      Hello Interval: 10, Hello Multiplier: 3, CSNP Interval: 10
    Number of active level-1 adjacencies: 1
    Level-2 Metric: 10, Priority: 64, Circuit ID: 0000.0000.000B.01
      Hello Interval: 10, Hello Multiplier: 3, CSNP Interval: 10
    Number of active level-2 adjacencies: 1
    Next IS-IS LAN Level-1 Hello in 4 seconds
    Next IS-IS LAN Level-2 Hello in 6 seconds
    LSP Interval: 33 Next IS-IS LAN Level-1 Hello in 4 seconds
    Next IS-IS LAN Level-2 Hello in 6 seconds
    LSP Interval: 33
  Restart Capable Neighbors: 2, In Start: 0, In Restart: 0
Dell#
```

Changing LSP Attributes

IS-IS routers flood link state PDUs (LSPs) to exchange routing information. LSP attributes include the generation interval, maximum transmission unit (MTU) or size, and the refresh interval.

You can modify the LSP attribute defaults, but it is not necessary.

To change the defaults, use any or all of the following commands.

- Set interval between LSP generation.

```
ROUTER ISIS mode
```

```
lsp-gen-interval [level-1 | level-2] seconds
```

– *seconds*: the range is from 0 to 120.

The default is **5 seconds**.

The default level is **Level 1**.

- Set the LSP size.

```
ROUTER ISIS mode
```

`lsp-mtu size`

- *size*: the range is from 128 to 9195.

The default is **1497**.

- Set the LSP refresh interval.

ROUTER ISIS mode

`lsp-refresh-interval seconds`

- *seconds*: the range is from 1 to 65535.

The default is **900 seconds**.

- Set the maximum time LSPs lifetime.

ROUTER ISIS mode

`max-lsp-lifetime seconds`

- *seconds*: the range is from 1 to 65535.

The default is **1200 seconds**.

Example of Viewing IS-IS Configuration (ROUTER ISIS Mode)

To view the configuration, use the `show config` command in ROUTER ISIS mode or the `show running-config isis` command in EXEC Privilege mode.

```
Dell#show running-config isis
!
router isis
  lsp-refresh-interval 902
  net 47.0005.0001.000C.000A.4321.00
  net 51.0005.0001.000C.000A.4321.00
Dell#
```

Configuring the IS-IS Metric Style

All IS-IS links or interfaces are associated with a cost that is used in the shortest path first (SPF) calculations. The possible cost varies depending on the metric style supported.

If you configure narrow, transition, or narrow transition metric style, the cost can be a number between 0 and 63. If you configure wide or wide transition metric style, the cost can be a number between 0 and 16,777,215. The system supports five different metric styles: narrow, wide, transition, narrow transition, and wide transition.

By default, the system generates and receives narrow metric values. Matrixes or costs higher than 63 are not supported. To accept or generate routes with a higher metric, you must change the metric style of the IS-IS process. For example, if you configure the metric as narrow, and a link state PDU (LSP) with wide metrics is received, the route is not installed.

The Dell Networking OS supports the following IS-IS metric styles.

Table 28. Metric Styles

Metric Style	Characteristics	Cost Range Supported on IS-IS Interfaces
narrow	Sends and accepts narrow or old TLVs (Type, Length, Value).	0 to 63
wide	Sends and accepts wide or new TLVs.	0 to 16777215
transition	Sends both wide (new) and narrow (old) TLVs.	0 to 63
narrow transition	Sends narrow (old) TLVs and accepts both narrow (old) and wide (new) TLVs.	0 to 63
wide transition	Sends wide (new) TLVs and accepts both narrow (old) and wide (new) TLVs.	0 to 16777215

To change the IS-IS metric style of the IS-IS process, use the following command.

- Set the metric style for the IS-IS process.

ROUTER ISIS mode

```
metric-style {narrow [transition] | transition | wide [transition]} [level-1
| level-2]
```

The default is **narrow**.

The default is Level 1 and Level 2 (**level-1-2**)

To view which metric types are generated and received, use the `show isis protocol` command in EXEC Privilege mode. The IS-IS matrixes settings are in bold.

Example of Viewing IS-IS Metric Types

```
Dell#show isis protocol
IS-IS Router: <Null Tag>
  System Id: EEEE.EEEE.EEEE IS-Type: level-1-2
  Manual area address(es):
    47.0004.004d.0001
  Routing for area address(es):
    21.2223.2425.2627.2829.3031.3233
    47.0004.004d.0001
  Interfaces supported by IS-IS:
    Vlan 2
    GigabitEthernet 4/22
    Loopback 0
  Redistributing:
  Distance: 115
  Generate narrow metrics: level-1-2
  Accept narrow metrics: level-1-2
  Generate wide metrics: none
  Accept wide metrics: none
Dell#
```

Configuring the IS-IS Cost

When you change from one IS-IS metric style to another, the IS-IS metric value could be affected. For each interface with IS-IS enabled, you can assign a cost or metric that is used in the link state calculation. To change the metric or cost of the interface, use the following commands.

- Assign an IS-IS metric.

INTERFACE mode

```
isis metric default-metric [level-1 | level-2]
```

- *default-metric*: the range is from 0 to 63 if the metric-style is narrow, narrow-transition, or transition.

The range is from 0 to 16777215 if the metric style is wide or wide transition.

The default is **10**.

- Assign a metric for an IPv6 link or interface.

INTERFACE mode

```
isis ipv6 metric default-metric [level-1 | level-2]
```

- *default-metric*: the range is from 0 to 63 for narrow and transition metric styles. The range is from 0 to 16777215 for wide metric styles.

The default is **10**.

The default level is **level-1**.

For more information about this command, refer to [Configuring the IS-IS Metric Style](#).

The following table describes the correct value range for the `isis metric` command.

Metric Sytle	Correct Value Range
wide	0 to 16777215
narrow	0 to 63
wide transition	0 to 16777215
narrow transition	0 to 63
transition	0 to 63

To view the interface's current metric, use the `show config` command in INTERFACE mode or the `show isis interface` command in EXEC Privilege mode.

Configuring the Distance of a Route

To configure the distance for a route, use the following command.

- Configure the distance for a route.

ROUTER ISIS mode

```
distance
```

Changing the IS-Type

To change the IS-type, use the following commands.

You can configure the system to act as a Level 1 router, a Level 1-2 router, or a Level 2 router.

To change the IS-type for the router, use the following commands.

- Configure IS-IS operating level for a router.

ROUTER ISIS mode

```
is-type {level-1 | level-1-2 | level-2-only}
```

Default is **level-1-2**.

- Change the IS-type for the IS-IS process.

ROUTER ISIS mode

```
is-type {level-1 | level-1-2 | level-2}
```

Example of the show isis database Command to View Level 1-2 Link State Databases

To view which IS-type is configured, use the `show isis protocol` command in EXEC Privilege mode.

The `show config` command in ROUTER ISIS mode displays only non-default information, so if you do not change the IS-type, the default value (**level-1-2**) is not displayed.

The default is Level 1-2 router. When the IS-type is Level 1-2, the software maintains two Link State databases, one for each level. To view the Link State databases, use the `show isis database` command.

```
Dell#show isis database
IS-IS Level-1 Link State Database
LSPID          LSP Seq Num  LSP Checksum LSP Holdtime ATT/P/OL
B233.00-00     0x00000003  0x07BF      1088         0/0/0
eljefe.00-00  * 0x00000009  0xF76A     1126         0/0/0
eljefe.01-00  * 0x00000001  0x68DF     1122         0/0/0
eljefe.02-00  * 0x00000001  0x2E7F     1113         0/0/0
Dell.00-00    0x00000002  0xD1A7     1102         0/0/0
IS-IS Level-2 Link State Database
LSPID          LSP Seq Num  LSP Checksum LSP Holdtime ATT/P/OL
B233.00-00     0x00000006  0xC38A     1124         0/0/0
eljefe.00-00  * 0x0000000D  0x51C6     1129         0/0/0
eljefe.01-00  * 0x00000001  0x68DF     1122         0/0/0
eljefe.02-00  * 0x00000001  0x2E7F     1113         0/0/0
Dell.00-00    0x00000004  0xCDA9     1107         0/0/0

Dell#
```

Controlling Routing Updates

To control the source of IS-IS route information, use the following command.

- Disable a specific interface from sending or receiving IS-IS routing information.

ROUTER ISIS mode

```
passive-interface interface
```

- For the Loopback interface on the RPM, enter the keyword `loopback` then a number from 0 to 16383.
- For a port channel, enter the keywords `port-channel` then a number from 1 to 255.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/ port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `FortyGigabitEthernet` then the slot/ port information.
- For a VLAN, enter the keyword `vlan` then a number from 1 to 4094.

Distribute Routes


Another method of controlling routing information is to filter the information through a prefix list.

Prefix lists are applied to incoming or outgoing routes and routes must meet the conditions of the prefix lists or the Dell Networking OS does not install the route in the routing table. The prefix lists are globally applied on all interfaces running IS-IS.

Configure the prefix list in PREFIX LIST mode prior to assigning it to the IS-IS process. For configuration information on prefix lists, refer to [Access Control Lists \(ACLs\)](#).

Applying IPv4 Routes

To apply prefix lists to incoming or outgoing IPv4 routes, use the following commands.

 **NOTE:** These commands apply to IPv4 IS-IS only. To apply prefix lists to IPv6 routes, use ADDRESS-FAMILY IPV6 mode, shown later.

- Apply a configured prefix list to all incoming IPv4 IS-IS routes.

ROUTER ISIS mode

```
distribute-list prefix-list-name in [interface]
```

- Enter the type of interface and slot/port information:
- For the Loopback interface on the RPM, enter the keyword `loopback` then a number from 0 to 16383.
- For a port channel, enter the keywords `port-channel` then a number from 1 to 255.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `FortyGigabitEthernet` then the slot/ port information.
- For a VLAN, enter the keyword `vlan` then a number from 1 to 4094.

- Apply a configured prefix list to all outgoing IPv4 IS-IS routes.

ROUTER ISIS mode

```
distribute-list prefix-list-name out [bgp as-number | connected | ospf process-id | rip | static]
```

You can configure one of the optional parameters:

- `connected`: for directly connected routes.
- `ospf process-id`: for OSPF routes only.
- `rip`: for RIP routes only.


- `static`: for user-configured routes.
- `bgp`: for BGP routes only.
- Deny RTM download for pre-existing redistributed IPv4 routes.

```
ROUTER ISIS mode

distribute-list redistributed-override in
```

Applying IPv6 Routes

To apply prefix lists to incoming or outgoing IPv6 routes, use the following commands.

 **NOTE:** These commands apply to IPv6 IS-IS only. To apply prefix lists to IPv4 routes, use ROUTER ISIS mode, previously shown.

- Apply a configured prefix list to all incoming IPv6 IS-IS routes.

```
ROUTER ISIS-AF IPV6 mode

distribute-list prefix-list-name in [interface]
```

Enter the type of interface and slot/port information:

- For the Loopback interface on the RPM, enter the keyword `loopback` then a number from 0 to 16383.
- For a port channel, enter the keywords `port-channel` then a number from 1 to 255.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `FortyGigabitEthernet` then the slot/port information.
- For a VLAN, enter the keyword `vlan` then a number from 1 to 4094.
- Apply a configured prefix list to all outgoing IPv6 IS-IS routes.

```
ROUTER ISIS-AF IPV6 mode
```

```
distribute-list prefix-list-name out [bgp as-number | connected | ospf
process-id | rip | static]
```

You can configure one of the optional parameters:


- `connected`: for directly connected routes.
- `ospf process-id`: for OSPF routes only.
- `rip`: for RIP routes only.
- `static`: for user-configured routes.
- `bgp`: for BGP routes only.
- Deny RTM download for pre-existing redistributed IPv6 routes.

```
ROUTER ISIS-AF IPV6 mode
```


```
distribute-list redistributed-override in
```

Redistributing IPv4 Routes

In addition to filtering routes, you can add routes from other routing instances or protocols to the IS-IS process. With the `redistribute` command syntax, you can include BGP, OSPF, RIP, static, or directly connected routes in the IS-IS process.

 **NOTE:** Do not route iBGP routes to IS-IS unless there are route-maps associated with the IS-IS redistribution.

To add routes from other routing instances or protocols, use the following commands.

 **NOTE:** These commands apply to IPv4 IS-IS only. To apply prefix lists to IPv6 routes, use ADDRESS-FAMILY IPV6 mode, shown later.

- Include BGP, directly connected, RIP, or user-configured (static) routes in IS-IS.

ROUTER ISIS mode

```
redistribute {bgp as-number | connected | rip | static} [level-1 level-1-2 |  
level-2] [metric metric-value] [metric-type {external | internal}] [route-map  
map-name]
```

Configure the following parameters:

- `level-1`, `level-1-2`, or `level-2`: assign all redistributed routes to a level. The default is **level-2**.
- `metric-value` the range is from 0 to 16777215. The default is **0**.
- `metric-type`: choose either external or internal. The default is **internal**.
- `map-name`: enter the name of a configured route map.

- Include specific OSPF routes in IS-IS.

ROUTER ISIS mode

```
redistribute ospf process-id [level-1 | level-1-2 | level-2] [metric value]  
[match external {1 | 2} | match internal] [metric-type {external | internal}]  
[route-map map-name]
```

Configure the following parameters:

- `process-id` the range is from 1 to 65535.
- `level-1`, `level-1-2`, or `level-2`: assign all redistributed routes to a level. The default is **level-2**.
- `metric value` the range is from 0 to 16777215. The default is **0**.
- `match external` the range is from 1 or 2.
- `match internal`
- `metric-type`: external or internal.
- `map-name`: enter the name of a configured route map.

Redistributing IPv6 Routes

To add routes from other routing instances or protocols, use the following commands.

 **NOTE:** These commands apply to IPv6 IS-IS only. To apply prefix lists to IPv4 routes, use the ROUTER ISIS mode previously shown.

- Include BGP, directly connected, RIP, or user-configured (static) routes in IS-IS.

ROUTER ISIS mode

```
redistribute {bgp as-number | connected | rip | static} [level-1 level-1-2 |
level-2] [metric metric-value] [metric-type {external | internal}] [route-map
map-name]
```

Configure the following parameters:

- `level-1`, `level-1-2`, or `level-2`: assign all redistributed routes to a level. The default is **level-2**.
 - `metric-value`: the range is from 0 to 16777215. The default is **0**.
 - `metric-type`: choose either external or internal. The default is **internal**.
 - `map-name`: enter the name of a configured route map.
- Include specific OSPF routes in IS-IS.ROUTER ISIS mode

```
redistribute ospf process-id [level-1| level-1-2 | level-2] [metric value]
[match external {1 | 2} | match internal] [metric-type {external | internal}]
[route-map map-name]
```

Configure the following parameters:

- `process-id`: the range is from 1 to 65535.
- `level-1`, `level-1-2`, or `level-2`: assign all redistributed routes to a level. The default is **level-2**.
- `metric value`: the range is from 0 to 16777215. The default is **0**.
- `match external`: the range is 1 or 2.
- `match internal`
- `metric-type`: external or internal.
- `map-name`: name of a configured route map.

To view the IS-IS configuration globally (including both IPv4 and IPv6 settings), use the `show running-config isis` command in EXEC Privilege mode. To view the current IPv4 IS-IS configuration, use the `show config` command in ROUTER ISIS mode. To view the current IPv6 IS-IS configuration, use the `show config` command in ROUTER ISIS-ADDRESS FAMILY IPV6 mode.

Configuring Authentication Passwords

You can assign an authentication password for routers in Level 1 and for routers in Level 2. Because Level 1 and Level 2 routers do not communicate with each other, you can assign different passwords for Level 1 routers and for Level 2 routers. However, if you want the routers in the level to communicate with each other, configure them with the same password.

To configure a simple text password, use the following commands.

- Configure authentication password for an area.

ROUTER ISIS mode

```
area-password [hmac-md5] password
```

The Dell Networking OS supports HMAC-MD5 authentication.

This password is inserted in Level 1 LSPs, Complete SNPs, and Partial SNPs.

- Set the authentication password for a routing domain.

ROUTER ISIS mode

```
domain-password [encryption-type | hmac-md5] password
```

The Dell Networking OS supports both DES and HMAC-MD5 authentication methods.

This password is inserted in Level 2 LSPs, Complete SNPs, and Partial SNPs.

To view the passwords, use the `show config` command in ROUTER ISIS mode or the `show running-config isis` command in EXEC Privilege mode.

To remove a password, use either the `no area-password` or `no domain-password` commands in ROUTER ISIS mode.

Setting the Overload Bit

Another use for the overload bit is to prevent other routers from using this router as an intermediate hop in their shortest path first (SPF) calculations. For example, if the IS-IS routing database is out of memory and cannot accept new LSPs, the system sets the overload bit and IS-IS traffic continues to transit the system.

To set or remove the overload bit manually, use the following commands.

- Set the overload bit in LSPs.

ROUTER ISIS mode

```
set-overload-bit
```

This setting prevents other routers from using it as an intermediate hop in their shortest path first (SPF) calculations.

- Remove the overload bit.

ROUTER ISIS mode

```
no set-overload-bit
```

Example of Viewing the Overload Bit Setting

When the bit is set, a 1 is placed in the *OL* column in the `show isis database` command output. The overload bit is set in both the Level-1 and Level-2 database because the IS type for the router is Level-1-2.

```
Dell#show isis database
IS-IS Level-1 Link State Database
LSPID          LSP Seq Num  LSP Checksum LSP Holdtime ATT/P/OL
B233.00-00     0x00000003   0x07BF       1074         0/0/0
eljefe.00-00 * 0x0000000A   0xF963       1196         0/0/1
eljefe.01-00 * 0x00000001   0x68DF       1108         0/0/0
eljefe.02-00 * 0x00000001   0x2E7F       1099         0/0/0
Dell.00-00     0x00000002   0xD1A7       1088         0/0/0
IS-IS Level-2 Link State Database
LSPID          LSP Seq Num  LSP Checksum LSP Holdtime ATT/P/OL
B233.00-00     0x00000006   0xC38A       1110         0/0/0
eljefe.00-00 * 0x0000000E   0x53BF       1196         0/0/1
eljefe.01-00 * 0x00000001   0x68DF       1108         0/0/0
eljefe.02-00 * 0x00000001   0x2E7F       1099         0/0/0
Dell.00-00     0x00000004   0xCDA9       1093         0/0/0
Dell#
```

Debugging IS-IS

To debug IS-IS processes, use the following commands.

- View all IS-IS information.

EXEC Privilege mode

```
debug isis
```

- View information on all adjacency-related activity (for example, hello packets that are sent and received).

EXEC Privilege mode

```
debug isis adj-packets [interface]
```

To view specific information, enter the following optional parameter:

- *interface*: Enter the type of interface and slot/port information to view IS-IS information on that interface only.

- View information about IS-IS local update packets.

EXEC Privilege mode

```
debug isis local-updates [interface]
```

To view specific information, enter the following optional parameter:

- *interface*: Enter the type of interface and slot/port information to view IS-IS information on that interface only.

- View IS-IS SNP packets, include CSNPs and PSNPs.

EXEC Privilege mode

```
debug isis snp-packets [interface]
```

To view specific information, enter the following optional parameter:

- *interface*: Enter the type of interface and slot/port information to view IS-IS information on that interface only.

- View the events that triggered IS-IS shortest path first (SPF) events for debugging purposes.

EXEC Privilege mode

```
debug isis spf-triggers
```

- View sent and received LSPs.

EXEC Privilege mode

```
debug isis update-packets [interface]
```

To view specific information, enter the following optional parameter:

- *interface*: Enter the type of interface and slot/port information to view IS-IS information on that interface only.

The system displays debug messages on the console. To view which debugging commands are enabled, use the `show debugging` command in EXEC Privilege mode.

To disable a specific debug command, enter the keyword `no` then the `debug` command. For example, to disable debugging of IS-IS updates, use the `no debug isis updates-packets` command.

To disable all IS-IS debugging, use the `no debug isis` command.

To disable all debugging, use the `undebug all` command.

IS-IS Metric Styles

The following sections provide additional information about the IS-IS metric styles.

- [Configuring the IS-IS Metric Style](#)
- [Configure Metric Values](#)

The Dell Networking OS supports the following IS-IS metric styles:

- narrow (supports only type, length, and value [TLV] up to 63)
- wide (supports TLV up to 16777215)
- transition (supports both narrow and wide and uses a TLV up to 63)
- narrow transition (accepts both narrow and wide and sends only narrow or old-style TLV)
- wide transition (accepts both narrow and wide and sends only wide or new-style TLV)

Configure Metric Values

For any level (Level-1, Level-2, or Level-1-2), the value range possible in the `isis metric` command in INTERFACE mode changes depending on the metric style.

The following describes the correct value range for the `isis metric` command.

Metric Style	Correct Value Range for the <code>isis metric</code> Command
wide	0 to 16777215
narrow	0 to 63
wide transition	0 to 16777215
narrow transition	0 to 63
transition	0 to 63

Maximum Values in the Routing Table

IS-IS metric styles support different cost ranges for the route. The cost range for the narrow metric style is 0 to 1023, while all other metric styles support a range of 0 to 0xFE000000.

Change the IS-IS Metric Style in One Level Only

By default, the IS-IS metric style is narrow. When you change from one IS-IS metric style to another, the IS-IS metric value (configured with the `isis metric` command) could be affected.

In the following scenarios, the IS-type is either Level-1 or Level-2 or Level-1-2 and the metric style changes.

Table 29. Metric Value When the Metric Style Changes

Beginning Metric Style	Final Metric Style	Resulting IS-IS Metric Value
wide	narrow	default value (10) if the original value is greater than 63. A message is sent to the console.
wide	transition	truncated value (the truncated value appears in the LSP only). The original <code>isis metric</code> value is displayed in the <code>show config</code> and <code>show running-config</code> commands and is used if you change back to transition metric style.
		 NOTE: A truncated value is a value that is higher than 63, but set back to 63 because the higher value is not supported.
wide	narrow transition	default value (10) if the original value is greater than 63. A message is sent to the console.
wide	wide transition	original value
narrow	wide	original value
narrow	transition	original value
narrow	narrow transition	original value
narrow	wide transition	original value
transition	wide	original value
transition	narrow	original value
transition	narrow	original value
transition	wide transition	original value
narrow transition	wide	original value
narrow transition	narrow	original value
narrow transition	wide transition	original value
narrow transition	transition	original value
wide transition	wide	original value
wide transition	narrow	default value (10) if the original value is greater than 63. A message is sent to the console.

Beginning Metric Style	Final Metric Style	Resulting IS-IS Metric Value
wide transition	narrow transition	default value (10) if the original value is greater than 63. A message is sent to the console.
wide transition	transition	truncated value (the truncated value appears in the LSP only). The original <code>isis metric</code> value is displayed in the <code>show config</code> and <code>show running-config</code> commands and is used if you change back to transition metric style.

Moving to transition and then to another metric style produces different results.

Table 30. Metric Value when the Metric Style Changes Multiple Times

Beginning Metric Style	Next Metric Style	Resulting Metric Value	Next Metric Style	Final Metric Value
wide	transition	truncated value	wide	original value is recovered
wide transition	transition	truncated value	wide transition	original value is recovered
wide	transition	truncated value	narrow	default value (10). A message is sent to the logging buffer
wide transition	transition	truncated value	narrow transition	default value (10). A message is sent to the logging buffer

Leaks from One Level to Another

In the following scenarios, each IS-IS level is configured with a different metric style.


Table 31. Metric Value with Different Levels Configured with Different Metric Styles

Level-1 Metric Style	Level-2 Metric Style	Resulting Metric Value
narrow	wide	original value
narrow	wide transition	original value
narrow	narrow transition	original value
narrow	transition	original value
wide	narrow	truncated value
wide	narrow transition	truncated value
wide	wide transition	original value
wide	transition	truncated value


Level-1 Metric Style	Level-2 Metric Style	Resulting Metric Value
narrow transition	wide	original value
narrow transition	narrow	original value
narrow transition	wide transition	original value
narrow transition	transition	original value
transition	wide	original value
transition	narrow	original value
transition	wide transition	original value
transition	narrow transition	original value
wide transition	wide	original value
wide transition	narrow	truncated value
wide transition	narrow transition	truncated value
wide transition	transition	truncated value

Sample Configurations

The following configurations are examples for enabling IPv6 IS-IS. These examples are not comprehensive directions. They are intended to give you some guidance with typical configurations.

 **NOTE:** Only one IS-IS process can run on the router, even if both IPv4 and IPv6 routing is being used.

You can copy and paste from these examples to your CLI. To support your own IP addresses, interfaces, names, and so on, be sure that you make the necessary changes.

 **NOTE:** Whenever you make IS-IS configuration changes, clear the IS-IS process (re-started) using the `clear isis` command. The `clear isis` command must include the tag for the ISIS process. The following example shows the response from the router:

```
Dell#clear isis *
% ISIS not enabled.
Dell#clear isis 9999 *
```

You can configure IPv6 IS-IS routes in one of the following three different methods:

- **Congruent Topology** — You *must* configure both IPv4 and IPv6 addresses on the interface. Enable the `ip router isis` and `ipv6 router isis` commands on the interface. Enable the `wide-metrics` parameter in router isis configuration mode.
- **Multi-topology** — You *must* configure the IPv6 address. Configuring the IPv4 address is optional. You *must* enable the `ipv6 router isis` command on the interface. If you configure IPv4, also enable the `router isis` command. In router isis configuration mode, enable `multi-topology` under `address-family ipv6 unicast`.
- **Multi-topology Transition** — You *must* configure the IPv6 address. Configuring the IPv4 address is optional. You *must* enable the `ipv6 router isis` command on the interface. If you configure IPv4, also enable the `ip router isis` command. In router isis configuration mode, enable `multi-topology transition` under `address-family ipv6 unicast`.




Figure 58. IPv6 IS-IS Sample Topography

IS-IS Sample Configuration — Router 1

IS-IS Sample Configuration — Router 2

IS-IS Sample Configuration — Router 3

The following is a sample configuration for enabling IPv6 IS-IS.

```
R1(conf)#interface Loopback 0
R1(conf-if-lo-0)#ip address 192.168.1.1/24
R1(conf-if-lo-0)#ipv6 address 2001:db8:9999:1::/48

R1(conf-if-lo-0)#ip router isis 9999
R1(conf-if-lo-0)#no shutdown
R1(conf-if-lo-0)#router isis 9999
R1(conf-router_isis)#is-type level-1
R1(conf-router_isis)#net FF.F101.0002.0C00.1111.00
R1(conf-router_isis)#ipv6 route 2001:db8:9999:2::/128 2001:db8:1021:2::
R1(conf)#ipv6 route 2001:db8:9999:3::/128 2001:db8:1022:3::
R1(conf)#ip route 192.168.1.2/32 10.0.12.2
R1(conf)#ip route 192.168.1.3/32 10.0.13.3

R1(conf)#interface GigabitEthernet 1/21
R1(conf-if-gi-1/21)#ip address 10.0.12.1/24
R1(conf-if-gi-1/21)#ipv6 address 2001:db8:1022:1::/48
R1(conf-if-gi-1/21)#isis circuit-type level-1
R1(conf-if-gi-1/21)#isis network point-to-point
R1(conf-if-gi-1/21)#ip router isis 9999
R1(conf-if-gi-1/21)#no shutdown

R1(conf-if-gi-1/21)#interface GigabitEthernet 1/34
R1(conf-if-gi-1/34)# ip address 10.0.13.1/24
R1(conf-if-gi-1/34)#ipv6 address 2001:db8:1021:1::/48
R1(conf-if-gi-1/34)#ip router isis 9999
R1(conf-if-gi-1/34)#no shutdown
R1(conf-if-gi-1/34)#end
```

```
R1#show ip route
Codes: C - connected, S - static, R - RIP,
       B - BGP, IN - internal BGP, EX - external BGP, LO - Locally Originated,
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,
       L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,
       > - non-active route, + - summary route
```

Gateway of last resort is not set

Destination	Gateway	Dist/Metric	Last Change
C 10.0.12.0/24	Direct, Gi 1/21	0/0	00:00:57
C 192.168.1.0/24	Direct, Lo 0	0/0	00:04:19
S 192.168.1.2/32	via 10.0.12.2, Gi 1/21	1/0	00:00:57

```
R1#show isis data
```

IS-IS Level-1 Link State Database

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
R1.00-00	* 0x0000000F	0x3A6C 1176		0/0/0
R1.02-00	* 0x00000002	0x90AC 1076		0/0/0
R1.03-00	* 0x00000002	0x67C3 1176		0/0/0
R2.00-00	0x0000000C	0x5418 1183		0/0/0
R2.00-00	0x00000009	0x1E39 1183		0/0/0
R2.03-00	0x00000002	0x589D 1180		0/0/0

```
R1#show isis neigh
```

System Id	Interface	State	Type	Priority	Uptime	Circuit Id
R2	Gi 1/21	Up	L1	64	00:02:28	R1.02
R2	Gi 1/34	Up	L1	64	00:00:42	R1.03

```
R1#
R2(conf)#interface Loopback 0
R2(conf-if-lo-0)#ip address 192.168.1.1/24
R2(conf-if-lo-0)#ipv6 address 2001:db8:9999:1::/48
R2(conf-if-lo-0)#ip router isis 9999
R2(conf-if-lo-0)#no shutdown
```

```
R2(conf-if-lo-0)#router isis 9999
R2(conf-router_isis)#int gi 2/11
R2(conf-if-gi-2/11)#ip address 10.0.12.2/24
R2(conf-if-gi-2/11)#ipv6 address 2001:db8:9999:2::/48
R2(conf-if-gi-2/11)#ip router isis 9999
R2(conf-if-gi-2/11)#isis network point-to-point
R2(conf-if-gi-2/11)#no shutdown
R2(conf-if-gi-2/11)#int gi 2/31
R2(conf-if-gi-2/31)#ip address 10.0.23.2/24
R2(conf-if-gi-2/31)#ipv6 address 2001:db8:1021:2::/48
R2(conf-if-gi-2/31)#ip router isis 9999
R2(conf-if-gi-2/31)#isis network point-to-point
R2(conf-if-gi-2/31)#no shutdown
R2(conf-if-gi-2/31)#router isis 9999
```

```
R2(conf-router_isis)#ipv6 route 2001:db8:9999:1::/128 2001:db8:1021:1::
R2(conf)#ipv6 route 2001:db8:9999:3::/128 2001:db8:1023:3::
R2(conf)#ip route 192.168.1.1/32 10.0.12.1
R2(conf)#ip route 192.168.1.3/32 10.0.23.3
R2(conf)#ex
```

```
R2#show ip route
```

```
Codes: C - connected, S - static, R - RIP,
       B - BGP, IN - internal BGP, EX - external BGP, LO - Locally Originated,
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
```

N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
 E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,
 L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,
 > - non-active route, + - summary route

Gateway of last resort is 172.21.212.1 to network 0.0.0.0

Destination	Gateway	Dist/Metric	Last Change
*S 0.0.0.0/0	via 172.21.212.1, V1	212 1/0	00:59:50
C 10.0.12.0/24	Direct, Gi 2/11	0/0	00:02:25
C 10.0.23.0/24	Direct, Gi 2/31	0/0	00:01:53
C 10.10.92.0/24	Direct, Po 4	0/0	6d9h
C 172.21.212.0/24	Direct, V1 212	0/0	2d20h
C 192.168.1.0/24	Direct, Lo 0	0/0	01:11:48
S 192.168.1.1/32	via 10.0.12.1, Gi 2/11	1/0	00:00:51
S 192.168.1.3/32	via 10.0.23.3, Gi 2/31	1/0	00:00:39

R2#show isis data

IS-IS Level-1 Link State Database

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
R2.00-00	* 0x0000000F	0x0174	1088	0/0/0

R2#show isis neigh

A2#show isis neigh

System Id	Interface	State	Type	Priority	Uptime	Circuit Id
R1	Gi 2/11	Up	L1	64	00:02:19	A102.02
R3	Gi 2/31	Up	L1	64	00:00:25	A121.03

R2#

R3(conf)#interface Loopback 0

R3(conf-if-lo-0)#ip address 192.168.1.3/24

R3(conf-if-lo-0)#ipv6 address 2001:db8:9999:3::/48

R3(conf-if-lo-0)#ip router isis 9999

R3(conf-if-lo-0)#no shutdown

R3(conf-if-lo-0)#router isis 9999

R3(conf-router_isis)#net FF.F101.0002.0C00.1133.00

R3(conf-router_isis)#ipv6 route 2001:db8:9999:1::/128 2001:db8:1022:1::

R3(conf)#ipv6 route 2001:db8:9999:2::/128 2001:db8:1023:2::

R3(conf)#ip route 192.168.1.1/32 10.0.13.1

R

3(conf)#interface GigabitEthernet 3/14

R3(conf-if-gi-3/14)#ip address 10.0.13.3/24

R3(conf-if-gi-3/14)#ipv6 address 2001:db8:1022:3::/48

R3(conf-if-gi-3/14)#ip router isis 9999

R3(conf-if-gi-3/14)#isis circuit-type level-1

R3(conf-if-gi-3/14)#isis network point-to-point

R3(conf-if-gi-3/14)#no shutdown

R3(conf-if-gi-3/14)#interface GigabitEthernet 3/21

R3(conf-if-gi-3/21)#ip address 10.0.23.3/24

R3(conf-if-gi-3/21)#ipv6 address 2001:db8:1023:3::/48

R3(conf-if-gi-3/21)#ip router isis 9999

R3(conf-if-gi-3/21)#isis circuit-type level-1

R3(conf-if-gi-3/21)#isis network point-to-point

R3(conf-if-gi-3/21)#no shutdown

R3(conf-if-gi-3/21)#end

R3#show ip route

Codes: C - connected, S - static, R - RIP,

B - BGP, IN - internal BGP, EX - external BGP, LO - Locally Originated,

O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,

N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,

E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,


```

L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,
> - non-active route, + - summary route
Gateway of last resort is not set
  Destination      Gateway             Dist/Metric  Last Change
-----
C 10.0.13.0/24    Direct, Gi 3/14    0/0         00:00:10
C 10.0.23.0/24    Direct, Gi 3/21    0/0         00:00:03
C 192.168.1.0/24  Direct, Lo 0       0/0         00:00:32
S 192.168.1.1/32  via 10.0.13.1, Gi 3/14  1/0         00:00:10
S 192.168.1.2/32  via 10.0.23.2, Gi 3/21  1/0         00:00:03

```

```
R2#show isis data
```

```

IS-IS Level-1 Link State Database
LSPID      LSP Seq Num  LSP Checksum LSP Holdtime ATT/P/OL
R1.00-00   0x0000000F   0x3A6C       1198         0/0/0
R1.03-00   0x00000001   0x69C2       1193         0/0/0
R2.00-00   * 0x00000007 0x51F6       1198         0/0/0
R2.03-00   * 0x00000001 0x5A9C       1200         0/0/0

```

```
IS-IS Level-2 Link State Database
```

```

LSPID      LSP Seq Num  LSP Checksum LSP Holdtime ATT/P/OL
R3.00-00   * 0x00000008 0xC09C       1199         0/0/0

```

```
R3#show isis neigh
```

```

System Id Interface State Type Priority Uptime  Circuit Id
R1          Gi 3/14   Init  L1    64      00:00:02 R1.03
R2          Gi 3/21   Up    L1    64      00:00:14 A101.03

```

Link Aggregation Control Protocol (LACP)

Link aggregation control protocol (LACP) is supported on the MXL switch platform.

Introduction to Dynamic LAGs and LACP

A link aggregation group (LAG), referred to as a *port channel*, can provide both load-sharing and port redundancy across line cards. You can enable LAGs as static or dynamic.

The benefits and constraints are basically the same, as described in *Port Channel Interfaces* in the [Interfaces](#) chapter.

The unique benefit of a dynamic LAG is that its ports can toggle between participating in the LAG or acting as dedicated ports, whereas ports in a static LAG must be removed from the LAG in order to act alone.

The Dell Networking operating system uses LACP to create dynamic LAGs. LACP provides a standardized means of exchanging information between two systems (also called Partner Systems) and automatically establishes the LAG between the systems. LACP permits the exchange of messages on a link to allow their LACP instances to:


- Reach an agreement on the identity of the LAG to which the link belongs.
- Move the link to that LAG.
- Enable the transmission and reception functions in an orderly manner.

The Dell Networking OS implementation of LACP is based on the standards specified in the IEEE 802.3: "Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications."

LACP functions by constantly exchanging custom MAC protocol data units (PDUs) across local area network (LAN) Ethernet links. The protocol packets are only exchanged between ports that are configured as LACP capable.

Important Points to Remember

- LACP allows you to add members to a port channel (LAG) as long as it has no static members. Conversely, if the LAG already contains a statically defined member (the `channel-member` command), the `port-channel mode` command is not permitted.
- A static LAG cannot be created if a dynamic LAG using the selected number exists.
- No dual membership in static and dynamic LAGs:
 - If a physical interface is a part of a static LAG, the `port-channel-protocol lacp` command is rejected on that interface.

- If a physical interface is a part of a dynamic LAG, it cannot be added as a member of a static LAG. The `channel-member gigabitethernet x/y` command is rejected in the static LAG interface for that physical interface.
- A dynamic LAG can be created with any type of configuration.
- There is a difference between the `shutdown` and `no interface port-channel` commands:
 - The `shutdown` command on LAG “xyz” disables the LAG and retains the user commands. However, the system does not allow the channel number “xyz” to be statically created.
 - The `no interface port-channel channel-number` command deletes the specified LAG, including a dynamically created LAG. This command removes all LACP-specific commands on the member interfaces. The interfaces are restored to a state that is ready to be configured.
-  **NOTE:** There is no configuration on the interface because that condition is required for an interface to be part of a LAG.
- You can configure link dampening on individual members of a LAG.
- You can configure a maximum of 128 port-channels with up to 16 members per channel.

LACP Modes

The Dell Networking OS provides three modes for configuration of LACP — Off, Active, and Passive.

- **Off** — In this state, an interface is not capable of being part of a dynamic LAG. LACP does not run on any port that is configured to be in this state.
- **Active** — In this state, the interface is said to be in the “active negotiating state.” LACP runs on any link that is configured to be in this state. A port in Active state also automatically initiates negotiations with other ports by initiating LACP packets.
- **Passive** — In this state, the interface is not in an active negotiating state, but LACP runs on the link. A port in Passive state also responds to negotiation requests (from ports in Active state). Ports in Passive state respond to LACP packets.

The Dell Networking OS supports LAGs in the following cases:

- A port in Active state can set up a port channel (LAG) with another port in Active state.
- A port in Active state can set up a LAG with another port in Passive state.
- A port in Passive state cannot set up a LAG with another port in Passive state.

Configuring LACP Commands

If you configure aggregated ports with compatible LACP modes (Off, Active, Passive), LACP can automatically link them, as defined in IEEE 802.3, Section 43.

To configure LACP, use the following commands.

- Configure the system priority.

CONFIGURATION mode

```
[no] lacp system-priority priority-value
```

The range is from 1 to 65535 (the higher the number, the lower the priority).

The default is **32768**.

- Enable or disable LACP on any LAN port.

INTERFACE mode

```
[no] port-channel-protocol lacp
```

The default is **LACP disabled**.

This command creates context.

- Configure LACP mode.

LACP mode

```
[no] port-channel number mode [active | passive | off]
```

– *number*: cannot statically contain any links.

The default is **LACP active**.

- Configure port priority.

LACP mode

```
[no] lacp port-priority priority-value
```

The range is from 1 to 65535 (the higher the number, the lower the priority).

The default is **32768**.

LACP Configuration Tasks

The following are LACP configuration tasks.

- [Creating a LAG](#)
- [Configuring the LAG Interfaces as Dynamic](#)
- [Setting the LACP Long Timeout](#)
- [Monitoring and Debugging LACP](#)
- [Configuring Shared LAG State Tracking](#)

Creating a LAG

To create a dynamic port channel (LAG), use the following command. First you define the LAG and then the LAG interfaces.

- Create a dynamic port channel (LAG).

CONFIGURATION mode

```
interface port-channel
```

- Create a dynamic port channel (LAG).

CONFIGURATION mode

```
switchport
```

Example of Configuring a LAG Interface

Example of the tagged Command

```
Dell(conf)#interface port-channel 32
Dell(conf-if-po-32)#no shutdown
Dell(conf-if-po-32)#switchport
```

The LAG is in the default VLAN. To place the LAG into a non-default VLAN, use the `tagged` command on the LAG.

```
Dell(conf)#interface vlan 10
Dell(conf-if-vl-10)#tagged port-channel 32
```

Configuring the LAG Interfaces as Dynamic

After creating a LAG, configure the dynamic LAG interfaces. To configure the dynamic LAG interfaces, use the following command.

- Configure the dynamic LAG interfaces.

CONFIGURATION mode

```
port-channel-protocol lacp
```

Example of the port-channel-protocol lacp Command


```
Dell(conf)#interface GigabitEthernet 3/15
Dell(conf-if-gi-3/15)#no shutdown
Dell(conf-if-gi-3/15)#port-channel-protocol lacp
Dell(conf-if-gi-3/15-lacp)#port-channel 32 mode active
...
Dell(conf)#interface GigabitEthernet 3/16
Dell(conf-if-gi-3/16)#no shutdown
Dell(conf-if-gi-3/16)#port-channel-protocol lacp
Dell(conf-if-gi-3/16-lacp)#port-channel 32 mode active
...
Dell(conf)#interface GigabitEthernet 4/15
Dell(conf-if-gi-4/15)#no shutdown
Dell(conf-if-gi-4/15)#port-channel-protocol lacp
Dell(conf-if-gi-4/15-lacp)#port-channel 32 mode active
...
Dell(conf)#interface GigabitEthernet 4/16
Dell(conf-if-gi-4/16)#no shutdown
Dell(conf-if-gi-4/16)#port-channel-protocol lacp
Dell(conf-if-gi-4/16-lacp)#port-channel 32 mode active
```

The `port-channel 32 mode active` command shown here may be successfully issued as long as there is no existing static channel-member configuration in LAG 32.

Setting the LACP Long Timeout

PDUs are exchanged between port channel (LAG) interfaces to maintain LACP sessions. PDUs are transmitted at either a slow or fast transmission rate, depending upon the LACP timeout value. The timeout value is the amount of time that a LAG interface waits for a PDU from the remote system before bringing the LACP session down. The default timeout value is **1 second**. You can configure the

default timeout value to be **30 seconds**. Invoking the longer timeout might prevent the LAG from flapping if the remote system is up but temporarily unable to transmit PDUs due to a system interruption.

 **NOTE:** The 30-second timeout is available for dynamic LAG interfaces only. You can enter the `lACP long-timeout` command for static LAGs, but it has no effect.

To configure LACP long timeout, use the following command.

- Set the LACP timeout value to 30 seconds.
CONFIG-INT-PO mode

```
lACP long-timeout
```

Example of the `lACP long-timeout` and `show lACP` Commands

```
Dell(conf)# interface port-channel 32
Dell(conf-if-po-32)#no shutdown
Dell(conf-if-po-32)#switchport
Dell(conf-if-po-32)#lACP long-timeout
Dell(conf-if-po-32)#end
Dell# show lACP 32
Port-channel 32 admin up, oper up, mode lACP
Actor System ID: Priority 32768, Address 0001.e800.a12b
Partner System ID: Priority 32768, Address 0001.e801.45a5
Actor Admin Key 1, Oper Key 1, Partner Oper Key 1
LACP LAG 1 is an aggregatable link
A - Active LACP, B - Passive LACP, C - Short Timeout, D - Long Timeout
E - Aggregatable Link, F - Individual Link, G - IN_SYNC, H - OUT_OF_SYNC
I - Collection enabled, J - Collection disabled, K - Distribution enabled L -
Distribution disabled,
M - Partner Defaulted, N - Partner Non-defaulted, O - Receiver is in expired
state,
P - Receiver is not in expired state
Port Gi 10/6 is enabled, LACP is enabled and mode is lACP
Actor Admin: State ADEHJLMP Key 1 Priority 128
```

To view the PDU exchanges and the timeout value, use the `debug lACP` command. For more information, refer to [Monitoring and Debugging LACP](#).

Monitoring and Debugging LACP

The system log (syslog) records faulty LACP actions.
To debug LACP, use the following command.

- Debug LACP, including configuration and events.
EXEC mode

```
[no] debug lACP [config | events | pdu [in | out | [interface [in | out]]]]
```

Shared LAG State Tracking

Shared LAG state tracking provides the flexibility to bring down a port channel (LAG) based on the operational state of another LAG.

At any time, only two LAGs can be a part of a group such that the fate (status) of one LAG depends on the other LAG.

As shown in the following illustration, the line-rate traffic from R1 destined for R4 follows the lowest-cost route via R2. Traffic is equally distributed between LAGs 1 and 2. If LAG 1 fails, all traffic from R1 to R4 flows across LAG 2 only. This condition over-subscribes the link and packets are dropped.




Figure 59. Shared LAG State Tracking

To avoid packet loss, redirect traffic through the next lowest-cost link (R3 to R4). The system has the ability to bring LAG 2 down if LAG 1 fails, so that traffic can be redirected. This redirection is what is meant by shared LAG state tracking. To achieve this functionality, you must group LAG 1 and LAG 2 into a single entity, called a failover group.

Configuring Shared LAG State Tracking

To configure shared LAG state tracking, you configure a failover group.

1. Enter port-channel failover group mode.
CONFIGURATION mode

`port-channel failover-group`
2. Create a failover group and specify the two port-channels that will be members of the group.
CONFIG-PO-FAILOVER-GRP mode

`group number port-channel number port-channel number`

Example of LAGs in the Same Failover Group

Example of Viewing the Failover Group Configuration

Example of Viewing Failover Group Member Status

In the following example, LAGs 1 and 2 have been placed into to the same failover group.

```
Dell#config
Dell(conf)#port-channel failover-group
Dell(conf-po-failover-grp)#group 1 port-channel 1 port-channel 2
```

To view the failover group configuration, use the `show running-configuration po-failover-group` command.

```
Dell#show running-config po-failover-group
!
port-channel failover-group
group 1 port-channel 1 port-channel 2
```

As shown in the following illustration, LAGs 1 and 2 are members of a failover group. LAG 1 fails and LAG 2 is brought down after the failure. This effect is logged by Message 1, in which a console message declares both LAGs down at the same time.

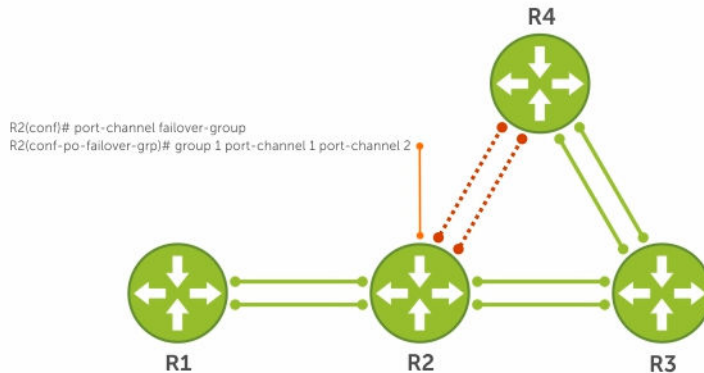


Figure 60. Configuring Shared LAG State Tracking

The following are shared LAG state tracking console messages:

- May 16 06:19:37: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down: Po 1
- May 16 06:19:37: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down: Po 2

To view the status of a failover group member, use the `show interface port-channel` command.

```
Dell#show interface port-channel 2
Port-channel 2 is up, line protocol is down (Failover-group 1 is down)
Hardware address is 00:01:e8:05:e8:4c, Current address is 00:01:e8:05:e8:4c
Interface index is 1107755010
Minimum number of links to bring Port-channel up is 1
Port-channel is part of failover-group 1
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 1000 Mbit
Members in this channel: Gi 1/17(U)
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:01:28
Queueing strategy: fifo
```

NOTE: The set of console messages shown above appear only if you configure shared LAG state tracking on that router (you can configure the feature on one or both sides of a link). For example, as previously shown, if you configured shared LAG state tracking on R2 only, no messages appear on R4 regarding the state of LAGs in a failover group.

Important Points about Shared LAG State Tracking

The following is more information about shared LAG state tracking.

- This feature is available for static and dynamic LAGs.
- Only a LAG can be a member of a failover group.
- You can configure shared LAG state tracking on one side of a link or on both sides.
- If a LAG that is part of a failover group is deleted, the failover group is deleted.
- If a LAG moves to the Down state due to this feature, its members may still be in the Up state.

LACP Basic Configuration Example

The screenshots in this section are based on the following example topology. Two routers are named ALPHA and BRAVO, and their hostname prompts reflect those names.




Figure 61. LACP Basic Configuration Example

Configure a LAG on ALPHA

The following example creates a LAG on ALPHA.

Example of Configuring a LAG

Example of Viewing a LAG Port Configuration

Summary of the LAG Configuration on Alpha

Summary of the LAG Configuration on Bravo

```
Alpha(conf)#interface port-channel 10
Alpha(conf-if-po-10)#no ip address
Alpha(conf-if-po-10)#switchport
Alpha(conf-if-po-10)#no shutdown
Alpha(conf-if-po-10)#show config
!
interface Port-channel 10
  no ip address
```

```
switchport
no shutdown
!
Alpha(conf-if-po-10)#
```

The following example inspects a LAG port configuration on ALPHA.

```
Alpha#sh int gig 2/31
GigabitEthernet 2/31 is up, line protocol is up
Port is part of Port-channel 10
Hardware is Force10Eth, address is 00:01:e8:06:95:c0
  Current address is 00:01:e8:06:95:c0
Interface Index is 109101113
Port will not be disabled on partial SFM failure
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 1000 Mbit, Mode full duplex, Slave
Flowcontrol rx on tx on
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:02:11
Queueing strategy: fifo
Input statistics:
  132 packets, 163668 bytes
  0 Vlans
  0 64-byte pkts, 12 over 64-byte pkts, 120 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  132 Multicasts, 0 Broadcasts
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics
  136 packets, 16718 bytes, 0 underruns
  0 64-byte pkts, 15 over 64-byte pkts, 121 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  136 Multicasts, 0 Broadcasts, 0 Unicasts
  0 Vlans, 0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec,0 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec,0 packets/sec, 0.00% of line-rate
  Time since last interface status change: 00:02:14
```

```

Alpha#sh int tengig 2/31
TenGigabitEthernet 2/31 is up, line protocol is up
Port is part of Port-channel 10
Hardware is Dell, address is 00:01:e8:06:95:c0
Current address is 00:01:e8:06:95:c0
Interface index is 109101113
Port will not be disabled on partial SFM failure
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 1000 Mbit, Mode full duplex, Slave
Flowcontrol rx on tx on
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:02:11
Queueing strategy: fifo
Input Statistics:
  132 packets, 16368 bytes
  0 Vlans
  0 64-byte pkts, 12 over 64-byte pkts, 120 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  132 Multicasts, 0 Broadcasts
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  136 packets, 16718 bytes, 0 underruns
  0 64-byte pkts, 15 over 64-byte pkts, 121 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  136 Multicasts, 0 Broadcasts, 0 Unicasts
  0 Vlans, 0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec,    0 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec,  0 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:02:14

```

Shows the status of this physical interface, and shows it is part of port channel 10.

Shows the speed of this physical interface. Also shows it is the slave of the TenGig link.

Figure 62. Inspecting the LAG Configuration

```
Alpha#show int port-channel 10
Port-channel 10 is up, line protocol is up
Created by LACP protocol
Hardware address is 00:01:e8:06:96:63, Current address is 00:01:e8:06:96:63
Interface index is 1107755018
Minimum number of links to bring Port-channel up is 1
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 3000 Mbit
Members in this channel: Gi 2/31(U) Gi 2/32(U) Gi 2/33(U)
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:04:09
Queueing strategy: fifo
Input Statistics:
  621 packets, 78732 bytes
  0 Vlans
  0 64-byte pkts, 18 over 64-byte pkts, 603 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  621 Multicasts, 0 Broadcasts
  0 runs, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  630 packets, 79284 bytes, 0 underruns
  0 64-byte pkts, 30 over 64-byte pkts, 600 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  630 Multicasts, 0 Broadcasts, 0 Unicasts
  0 Vlans, 0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec,      2 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec,    2 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:03:38
```

Shows the speed of this physical interface. Also shows it is the slave of the TenGig link.

Confirms the number of links to bring up the LAG and that this is a switch port instead of a router port.

Confirms the number of links to bring up the LAG and that this is a switch port instead of a router port.

Figure 63. Inspecting Configuration of LAG 10 on ALPHA

```

Alpha#sho lacp 10
Port-channel 10 admin up, oper up, mode lacp
Actor System ID: Priority 32768, Address 0001.e806.953e
Partner System ID: Priority 32768, Address 0001.e809.c24a
Actor Admin Key 10, Oper Key 10, Partner Oper Key 10
LACP LAG 10 is an aggregatable link

A - Active LACP, B - Passive LACP, C - Short Timeout, D - Long Timeout
E - Aggregatable Link, F - Individual Link, G - IN_SYNC, H - OUT_OF_SYNC
I - Collection enabled, J - Collection disabled, K - Distribution enabled
L - Distribution disabled, M - Partner Defaulted, N - Partner Non-defaulted,
O - Receiver is in expired state, P - Receiver is not in expired state

Port Gi 2/31 is enabled, LACP is enabled and mode is lacp
Actor Admin: State ACEHJLMP Key 10 Priority 32768
Oper: State ACEGIKNP Key 10 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
Oper: State ACEGIKNP Key 10 Priority 32768

Port Gi 2/32 is enabled, LACP is enabled and mode is lacp
Actor Admin: State ACEHJLMP Key 10 Priority 32768
Oper: State ACEGIKNP Key 10 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
Oper: State ACEGIKNP Key 10 Priority 32768

Port Gi 2/33 is enabled, LACP is enabled and mode is lacp
Actor Admin: State ACEHJLMP Key 10 Priority 32768
Oper: State ACEGIKNP Key 10 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
Oper: State ACEGIKNP Key 10 Priority 32768
Alpha#

```

Figure 64. Verifying LAG 10 Status on ALPHA Using the show lacp Command

```

Alpha(conf-if-po-10)#int gig 2/31
Alpha(conf-if-gi-2/31)#no ip address
Alpha(conf-if-gi-2/31)#no switchport
Alpha(conf-if-gi-2/31)#shutdown
Alpha(conf-if-gi-2/31)#port-channel-protocol lacp
Alpha(conf-if-gi-2/31-lacp)#port-channel 10 mode active
Alpha(conf-if-gi-2/31-lacp)#no shut
Alpha(conf-if-gi-2/31)#show config

!
interface GigabitEthernet 2/31
  no ip address
!
  port-channel-protocol LACP
  port-channel 10 mode active
  no shutdown
!
Alpha(conf-if-gi-2/31)#

interface Port-channel 10
no ip address
switchport
no shutdown

```

```

interface GigabitEthernet 2/31
no ip address

Bravo(conf-if-gi-3/21)#int port-channel 10
Bravo(conf-if-po-10)#no ip add
Bravo(conf-if-po-10)#switch
Bravo(conf-if-po-10)#no shut
Bravo(conf-if-po-10)#show config
!
interface Port-channel 10
no ip address
switchport
no shutdown
!
Bravo(conf-if-po-10)#exit

Bravo(conf)#int gig 3/21
Bravo(conf)#no ip address
Bravo(conf)#no switchport
Bravo(conf)#shutdown
Bravo(conf-if-gi-3/21)#port-channel-protocol lacp
Bravo(conf-if-gi-3/21-lacp)#port-channel 10 mode active
Bravo(conf-if-gi-3/21-lacp)#no shut
Bravo(conf-if-gi-3/21)#end

!
interface GigabitEthernet 3/21
no ip address
!
port-channel-protocol LACP
port-channel 10 mode active
no shutdown
Bravo(conf-if-gi-3/21)#end

int port-channel 10
no ip address
switchport
no shutdown
show config

int gig 3/21
no ip address

```

```
Bravo#show int tengig 3/21
GigabitEthernet 3/21 is up, line protocol is up
Port is part of Port-channel 10
Hardware is Dell, address is 00:01:e8:09:c3:82
Current address is 00:01:e8:09:c3:82
Interface index is 140034106
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 1000 Mbit, Mode full duplex, Master
Flowcontrol rx on tx on
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:15:05
Queueing strategy: fifo
Input Statistics:
  708 packets, 89934 bytes
  0 Vlans
  0 64-byte pkts, 15 over 64-byte pkts, 693 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  708 Multicasts, 0 Broadcasts
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  705 packets, 89712 bytes, 0 underruns
  0 64-byte pkts, 12 over 64-byte pkts, 693 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  705 Multicasts, 0 Broadcasts, 0 Unicasts
  0 Vlans, 0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec,      0 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec,    0 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:12:39
```

Shows the status of this interface. Also shows it is part of LAG 10.

Shows that this is a Layer 2 port.

Shows the speed of this physical interface. Also shows it is the Master of the TenGig link.

Figure 65. Inspecting a LAG Port on BRAVO Using the show interface Command

```

Dell#sh int port 10
Port-channel 10 is up, line protocol is up
Created by LACP protocol
Hardware address is 00:01:e8:09:c4:ef, Current address is 00:01:e8:09:c4:ef
Interface index is 1107755018
Minimum number of links to bring Port-channel up is 1
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 3000 Mbit
Members in this channel: Gi 3/21(U) Gi 3/22(U) Gi 3/23(U)
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:13:07
Queueing strategy: fifo
Input Statistics:
  2189 packets, 278744 bytes
  0 Vlans
  0 64-byte pkts, 32 over 64-byte pkts, 2157 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  2189 Multicasts, 0 Broadcasts
  0 runs, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  2173 packets, 277350 bytes, 0 underruns
  0 64-byte pkts, 19 over 64-byte pkts, 2154 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  2173 Multicasts, 0 Broadcasts, 0 Unicasts
  0 Vlans, 0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec,      2 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec,    2 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:13:00

```

Indicates the MAC address assigned to the LAG. This does NOT match any of the physical interface MAC addresses.

Confirms the number of links to bring up the LAG and that this is a switch port instead of a router port.

Confirms the total bandwidth for this LAG and which interfaces are active.

Figure 66. Inspecting LAG 10 Using the show interfaces port-channel Command


```

Dell#sh int port 10
Port-channel 10 is up, line protocol is up
Created by LACP protocol
Hardware address is 00:01:e8:09:c4:ef, Current address is 00:01:e8:09:c4:ef
Interface index is 1107755018
Minimum number of links to bring Port-channel up is 1
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 3000 Mbit
Members in this channel: Gi 3/21(U) Gi 3/22(U) Gi 3/23(U)
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:13:07
Queueing strategy: fifo
Input Statistics:
  2189 packets, 278744 bytes
  0 Vlans
  0 64-byte pkts, 32 over 64-byte pkts, 2157 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  2189 Multicasts, 0 Broadcasts
  0 runs, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  2173 packets, 277350 bytes, 0 underruns
  0 64-byte pkts, 19 over 64-byte pkts, 2154 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  2173 Multicasts, 0 Broadcasts, 0 Unicasts
  0 Vlans, 0 throttles, 0 discarded, 0 collisions, 0 wredrops
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec,      2 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec,    2 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:13:00

```

Indicates the MAC address assigned to the LAG. This does NOT match any of the physical interface MAC addresses.

Confirms the number of links to bring up the LAG and that this is a switch port instead of a router port.

Confirms the total bandwidth for this LAG and which interfaces are active.

Figure 67. Inspecting the LAG Status Using the show lacp command

The point-to-point protocol (PPP) is a connection-oriented protocol that enables layer two links over various different physical layer connections. It is supported on both synchronous and asynchronous lines, and can operate in Half-Duplex or Full-Duplex mode. It was designed to carry IP traffic but is general enough to allow any type of network layer datagram to be sent over a PPP connection. As its name implies, it is for point-to-point connections between exactly two devices, and assumes that frames are sent and received in the same order.

Layer 2

Layer 2 features are supported on the MXL switch platform.

Manage the MAC Address Table

The Dell Networking OS provides the following management activities for the MAC address table.

- [Clearing the MAC Address Table](#)
- [Setting the Aging Time for Dynamic Entries](#)
- [Configuring a Static MAC Address](#)
- [Displaying the MAC Address Table](#)

Clearing the MAC Address Table

You may clear the MAC address table of dynamic entries. To clear a MAC address table, use the following command.

- Clear a MAC address table of dynamic entries.
EXEC Privilege mode

```
clear mac-address-table dynamic {address | all | interface | vlan}
```

- *address*: deletes the specified entry.
- *all*: deletes all dynamic entries.
- *interface*: deletes all entries for the specified interface.
- *vlan*: deletes all entries for the specified VLAN.

Setting the Aging Time for Dynamic Entries

Learned MAC addresses are entered in the table as dynamic entries, which means that they are subject to aging.

For any dynamic entry, if no packet arrives on the switch with the MAC address as the source or destination address within the timer period, the address is removed from the table. The default aging time is **1800 seconds**.

To disable a MAC address and specify an aging time, use the following commands.

- Disable MAC address aging for all dynamic entries.
CONFIGURATION mode

```
mac-address-table aging-time 0
```

- Specify an aging time.

CONFIGURATION mode

```
mac-address-table aging-time seconds
```

The range is from 10 to 1000000.

Dell Networking OS Behavior: The time elapsed before the configured MAC aging time expires is not precisely as configured. For example, the VLAN configuration `mac-address-table aging-time 1`, does not remove dynamic entries from the CAM after precisely 1 second. The actual minimum aging time for entries is approximately 5 seconds because this is the default MAC address table scanning interval. Therefore, MAC aging configurations of less than 5 seconds, as in this example, might be ineffective. Configuring `mac-address-table station-move time-interval 500` solves this limitation. Reducing the scanning interval to the minimum (500 milliseconds), increases the detection speed, which results in the system clearing entries closer to the actual desired aging time.

Configuring a Static MAC Address

A static entry is one that is not subject to aging. Enter static entries manually. To create a static MAC address entry, use the following command.

- Create a static MAC address entry in the MAC address table.

CONFIGURATION mode

```
mac-address-table static
```

Displaying the MAC Address Table

To display the MAC address table, use the following command.

- Display the contents of the MAC address table.

EXEC Privilege mode

```
show mac-address-table [address | aging-time [vlan vlan-id] | count | dynamic  
| interface | static | vlan]
```

- `address`: displays the specified entry.
- `aging-time`: displays the configured aging-time.
- `count`: displays the number of dynamic and static entries for all VLANs, and the total number of entries.
- `dynamic`: displays only dynamic entries.
- `interface`: displays only entries for the specified interface.
- `static`: displays only static entries.
- `vlan`: displays only entries for the specified VLAN.

MAC Learning Limit

MAC address learning limit is a method of port security on Layer 2 port-channel and physical interfaces, and VLANs. It allows you to set an upper limit on the number of MAC addresses that learned on an

interface/VLAN. After the limit is reached, the system drops all traffic from a device with an unlearned MAC address.

This section describes the following:

- [mac learning-limit Dynamic](#)
- [mac learning-limit station-move](#)
- [Learning Limit Violation Actions](#)
- [Setting Station Move Violation Actions](#)
- [Recovering from Learning Limit and Station Move Violations](#)

Dell Networking OS Behavior: When configuring the MAC learning limit on a port, the configuration is accepted (becomes part of `running-config` and `show mac learning-limit interface`) before the system verifies that sufficient CAM space exists. If the CAM check fails, a message displays:

```
%E90MH:5 %ACL_AGENT-2-ACL_AGENT_LIST_ERROR: Unable to apply access-list Mac-Limit on GigabitEthernet 5/84
```

In this case, the configuration is still present in the `running-config` and `show` output. Remove the configuration before re-applying a MAC learning limit with a lower value. Also, ensure that you can view the Syslog messages on your session.

Setting the MAC Learning Limit

To set a MAC learning limit on an interface, use the following command.

- Specify the number of MAC addresses that the system can learn off a Layer 2 interface.

INTERFACE mode

```
mac learning-limit address_limit
```

Three options are available with the `mac learning-limit` command:

- ✎ **NOTE:** An SNMP trap is available for `mac learning-limit station-move`. No other SNMP traps are available for MAC Learning Limit, including limit violations.

mac learning-limit Dynamic

The MAC address table is stored on the Layer 2 forwarding information base (FIB) region of the CAM.

The Layer 2 FIB region allocates space for static MAC address entries and dynamic MAC address entries. When you enable MAC learning limit, entries created on this port are static by default. When you configure the `dynamic` option, learned MAC addresses are stored in the dynamic region and are subject to aging. Entries created before this option is set are not affected.

Dell Networking OS Behavior: If you do not configure the `dynamic` option, the MXL switch does not detect station moves in which a MAC address learned off of a MAC-limited port is learned on another port on the same stack unit or different stack.

mac learning-limit station-move

The `station-move` option, allows a MAC address already in the table to be learned off of another interface.

For example, if you disconnect a network device from one interface and reconnect it to another interface, the MAC address is learned on the new interface. When the system detects this “station move,” the system clears the entry learned on the original interface and installs a new entry on the new interface.

Learning Limit Violation Actions

To configure the system to take an action when the MAC learning limit is reached on an interface and a new address is received using one of the following options with the `mac learning-limit` command, use the following commands.

- Generate a system log message when the MAC learning limit is exceeded.
INTERFACE mode

`learn-limit-violation log`
- Shut down the interface and generate a system log message when the MAC learning limit is exceeded.
INTERFACE mode

`learn-limit-violation shutdown`

Setting Station Move Violation Actions

`no-station-move` is the default behavior. You can configure the system to take an action if a station move occurs using one of the following options with the `mac learning-limit` command.

To display a list of interfaces configured with MAC learning limit or station move violation actions, use the following commands.

- Generate a system log message indicating a station move.
INTERFACE mode

`station-move-violation log`
- Shut down the first port to learn the MAC address.
INTERFACE mode


`station-move-violation shutdown-original`
- Shut down the second port to learn the MAC address.
INTERFACE mode

`station-move-violation shutdown-offending`
- Shut down both the first and second port to learn the MAC address.
INTERFACE mode

`station-move-violation shutdown-both`


- Display a list of all of the interfaces configured with MAC learning limit or station move violation.
CONFIGURATION mode

```
show mac learning-limit violate-action
```

 **NOTE:** When the MAC learning limit (MLL) is configured as `no-station-move`, the MLL will be processed as static entries internally. For static entries, the MAC address will be installed in all port-pipes, irrespective of the VLAN membership.

Recovering from Learning Limit and Station Move Violations

After a learning-limit or station-move violation shuts down an interface, you must manually reset it. To reset the learning limit, use the following commands.

 **NOTE:** Alternatively, you can reset the interface by shutting it down using the `shutdown` command and then re-enabling it using the `no shutdown` command.

- Reset interfaces in the ERR_Disabled state caused by a learning limit violation or station move violation.

EXEC Privilege mode

```
mac learning-limit reset
```

- Reset interfaces in the ERR_Disabled state caused by a learning limit violation.

EXEC Privilege mode

```
mac learning-limit reset learn-limit-violation [interface | all]
```

- Reset interfaces in the ERR_Disabled state caused by a station move violation.

EXEC Privilege mode

```
mac learning-limit reset station-move-violation [interface | all]
```

NIC Teaming

NIC teaming is a feature that allows multiple network interface cards in a server to be represented by one MAC address and one IP address in order to provide transparent redundancy, balancing, and to fully utilize network adapter resources.

The following illustration shows a topology where two NICs have been teamed together. In this case, if the primary NIC fails, traffic switches to the secondary NIC because they are represented by the same set of addresses.

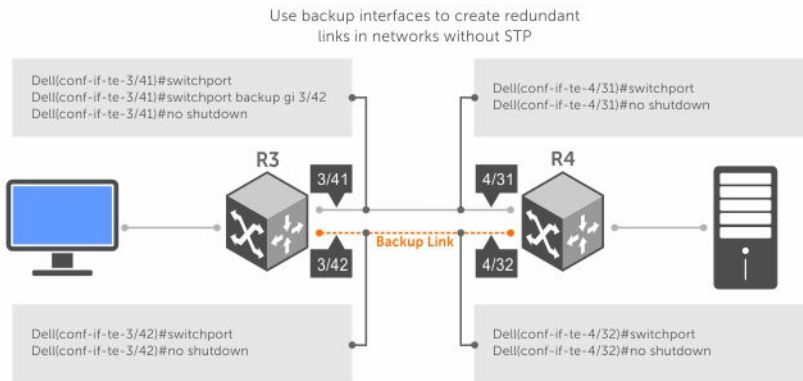



Figure 68. Redundant NICs with NIC Teaming

When you use NIC teaming, consider that the server MAC address is originally learned on Port 0/1 of the switch (shown in the following) and Port 0/5 is the failover port. When the NIC fails, the system automatically sends an ARP request for the gateway or host NIC to resolve the ARP and refresh the egress interface. When the ARP is resolved, the same MAC address is learned on the same port where the ARP is resolved (in the previous example, this location is Port 0/5 of the switch). To ensure that the MAC address is disassociated with one port and re-associated with another port in the ARP table, configure the `mac-address-table station-move refresh-arp` command on the Dell Networking switch at the time that NIC teaming is being configured on the server.

NOTE: If you do not configure the `mac-address-table station-move refresh-arp` command, traffic continues to be forwarded to the failed NIC until the ARP entry on the switch times out.

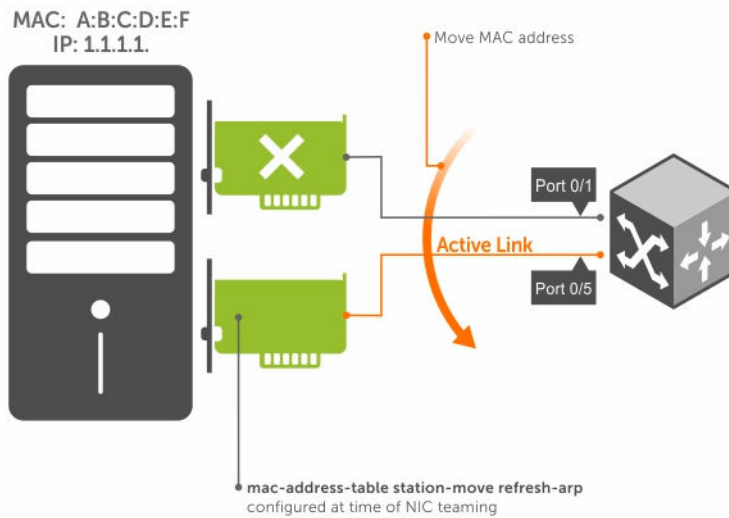


Figure 69. Configuring the mac-address-table station-move refresh-arp Command

MAC Move Optimization

MAC move optimization is supported only on the E-Series platform.

Station-move detection takes 5000ms because this is the interval at which the detection algorithm runs.

The `threshold` option is the number of times a station move must be detected in a single interval in order to trigger a system log message. For example, if you configure `mac-address-table station-move threshold 2 time-interval 5000`, and four station moves occur in 5000ms, then two log messages are generated.

Link Layer Discovery Protocol (LLDP)

The link layer discovery protocol (LLDP) is supported on the MXL switch platform.

802.1AB (LLDP) Overview

LLDP — defined by IEEE 802.1AB — is a protocol that enables a local area network (LAN) device to advertise its configuration and receive configuration information from adjacent LLDP-enabled LAN infrastructure devices.

The collected information is stored in a management information base (MIB) on each device, and is accessible via simple network management protocol (SNMP).

Protocol Data Units

Configuration information is exchanged in the form of Type, Length, Value (TLV) segments.

- Type — The kind of information included in the TLV.
- Length — The value, in octets, of the TLV after the Length field.
- Value — The configuration information that the agent is advertising.

The chassis ID TLV is shown in the following illustration.

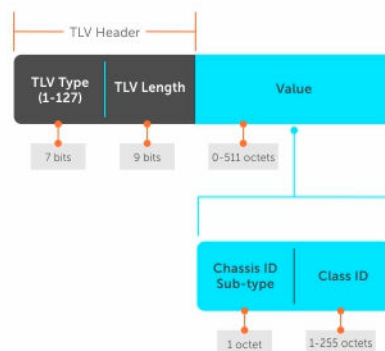


Figure 70. Type, Length, Value (TLV) Segment

TLVs are encapsulated in a frame called an LLDP data unit (LLDPDU) (shown in the following table), which is transmitted from one LLDP-enabled device to its LLDP-enabled neighbors. LLDP is a one-way protocol. LLDP-enabled devices (LLDP agents) can transmit and/or receive advertisements, but they cannot solicit and do not respond to advertisements.

There are five types of TLVs. All types are mandatory in the construction of an LLDPDU except Optional TLVs. You can configure the inclusion of individual Optional TLVs.

Table 32. Type, Length, Value (TLV) Types

Type	TLV	Description
0	End of LLDPDU	Marks the end of an LLDPDU.
1	Chassis ID	An administratively assigned name that identifies the LLDP agent.
2	Port ID	An administratively assigned name that identifies a port through which TLVs are sent and received.
3	Time to Live	A value that tells the receiving agent how long the information contained in the TLV Value field is valid.
—	Optional	Includes sub-types of TLVs that advertise specific configuration information. These sub-types are Management TLVs, IEEE 802.1, IEEE 802.3, and TIA-1057 Organizationally Specific TLVs.




Figure 71. LLDPDU Frame

Optional TLVs

The Dell Networking operating system supports these optional TLVs: management TLVs, IEEE 802.1 and 802.3 organizationally specific TLVs, and TIA-1057 organizationally specific TLVs.

Management TLVs

A management TLV is an optional TLVs sub-type. This kind of TLV contains essential management information about the sender.

Organizationally Specific TLVs

A professional organization or a vendor can define organizationally specific TLVs. They have two mandatory fields (as shown in the following illustration) in addition to the basic TLV fields.

- Organizationally Unique Identifier (OUI)—a unique number assigned by the IEEE to an organization or vendor.
- OUI Sub-type—These sub-types indicate the kind of information in the following data field. The sub-types are determined by the owner of the OUI.




Figure 72. Organizationally Specific TLV

IEEE Organizationally Specific TLVs

Eight TLV types have been defined by the IEEE 802.1 and 802.3 working groups as a basic part of LLDP; the IEEE OUI is 00-80-C2. You can configure the Dell Networking system to advertise any or all of these TLVs.

Table 33. Optional TLV Types

Type	TLV	Description
Optional TLVs		
4	Port description	A user-defined alphanumeric string that describes the port. The Dell Networking OS does not currently support this TLV.
5	System name	A user-defined alphanumeric string that identifies the system.
6	System description	A user-defined alphanumeric string that identifies the system.
7	System capabilities	Identifies the chassis as one or more of the following: repeater, bridge, WLAN Access Point, Router, Telephone, DOCSIS cable device, end station only, or other.
8	Management address	Indicates the network address of the management interface. The

Type	TLV	Description
		Dell Networking OS does not currently support this TLV.
IEEE 802.1 Organizationally Specific TLVs		
127	Port-VLAN ID	On Dell Networking systems, indicates the untagged VLAN to which a port belongs.
127	Port and Protocol VLAN ID	On Dell Networking systems, indicates the tagged VLAN to which a port belongs (and the untagged VLAN to which a port belongs if the port is in Hybrid mode).
127	VLAN Name	Indicates the user-defined alphanumeric string that identifies the VLAN.
127	Protocol Identity	Indicates the protocols that the port can process. The Dell Networking OS does not currently support this TLV.
IEEE 802.3 Organizationally Specific TLVs		
127	MAC/PHY Configuration/Status	Indicates the capability and current setting of the duplex status and bit rate, and whether the current settings are the result of auto-negotiation. This TLV is not available in the the Dell Networking OS implementation of LLDP, but is available and mandatory (non-configurable) in the LLDP-MED implementation.
127	Power via MDI	Dell Networking supports the LLDP-MED protocol, which recommends that Power via MDI TLV be not implemented, and therefore Dell Networking implements Extended Power via MDI TLV only.
127	Link Aggregation	Indicates whether the link is capable of being aggregated, whether it is currently in a LAG, and the port identification of the LAG. The Dell Networking OS
536		Link Layer Discovery Protocol (LLDP)

Type	TLV	Description
127	Maximum Frame Size	does not currently support this TLV. Indicates the maximum frame size capability of the MAC and PHY.

TIA-1057 (LLDP-MED) Overview

Link layer discovery protocol – media endpoint discovery (LLDP-MED) as defined by ANSI/ TIA-1057– provides additional organizationally specific TLVs so that endpoint devices and network connectivity devices can advertise their characteristics and configuration information; the OUI for the Telecommunications Industry Association (TIA) is 00-12-BB.

- **LLDP-MED Endpoint Device** – any device that is on an IEEE 802 LAN network edge can communicate using IP and uses the LLDP-MED framework.
- **LLDP-MED Network Connectivity Device** – any device that provides access to an IEEE 802 LAN to an LLDP-MED endpoint device and supports IEEE 802.1AB (LLDP) and TIA-1057 (LLDP-MED). The Dell Networking system is an LLDP-MED network connectivity device.

Regarding connected endpoint devices, LLDP-MED provides network connectivity devices with the ability to:

- manage inventory
- manage Power over Ethernet (PoE)
- identify physical location
- identify network policy

LLDP-MED is designed for, but not limited to, VoIP endpoints.

TIA Organizationally Specific TLVs

The Dell Networking system is an LLDP-MED Network Connectivity Device (Device Type 4).

Network connectivity devices are responsible for:

- transmitting an LLDP-MED capability TLV to endpoint devices
- storing the information that endpoint devices advertise

The following table describes the five types of TIA-1057 Organizationally Specific TLVs.

Table 34. TIA-1057 (LLDP-MED) Organizationally Specific TLVs

Type	SubType	TLV	Description
127	1	LLDP-MED Capabilities	Indicates: <ul style="list-style-type: none"> • whether the transmitting device supports LLDP-MED • what LLDP-MED TLVs it supports

Type	SubType	TLV	Description
			<ul style="list-style-type: none"> LLDP device class
127	2	Network Policy	Indicates the application type, VLAN ID, Layer 2 Priority, and DSCP value.
127	3	Location Identification	<p>Indicates that the physical location of the device expressed in one of three possible formats:</p> <ul style="list-style-type: none"> Coordinate Based LCI Civic Address LCI Emergency Call Services ELIN
127	4	Location Identification	Indicates power requirements, priority, and power status.
Inventory Management TLVs	Implementation of this set of TLVs is optional in LLDP-MED devices. None or all TLVs must be supported. The Dell Networking OS does not currently support these TLVs.		
127	5	Inventory – Hardware Revision	Indicates the hardware revision of the LLDP-MED device.
127	6	Inventory – Firmware Revision	Indicates the firmware revision of the LLDP-MED device.
127	7	Inventory – Software Revision	Indicates the software revision of the LLDP-MED device.
127	8	Inventory – Serial Number	Indicates the device serial number of the LLDP-MED device.
127	9	Inventory – Manufacturer Name	Indicates the manufacturer of the LLDP-MED device.
127	10	Inventory – Model Name	Indicates the model of the LLDP-MED device.

Type	SubType	TLV	Description
127	11	Inventory – Asset ID	Indicates a user specified device number to manage inventory.
127	12–255	Reserved	–

LLDP-MED Capabilities TLV

The LLDP-MED capabilities TLV communicates the types of TLVs that the endpoint device and the network connectivity device support. LLDP-MED network connectivity devices must transmit the Network Policies TLV.

- The value of the LLDP-MED capabilities field in the TLV is a 2–octet bitmap, each bit represents an LLDP-MED capability (as shown in the following table).
- The possible values of the LLDP-MED device type are shown in the following. The Dell Networking system is a network connectivity device, which is Type 4.

When you enable LLDP-MED in Dell Networking OS (using the `advertise med` command), the system begins transmitting this TLV.



Figure 73. LLDP-MED Capabilities TLV

Table 35. Dell Networking OS LLDP-MED Capabilities

Bit Position	TLV	Dell Networking OS Support
0	LLDP-MED Capabilities	Yes
1	Network Policy	Yes
2	Location Identification	Yes
3	Extended Power via MDI-PSE	Yes
4	Extended Power via MDI-PD	No
5	Inventory	No
6–15	reserved	No

Table 36. LLDP-MED Device Types

Value	Device Type
0	Type Not Defined
1	Endpoint Class 1
2	Endpoint Class 2
3	Endpoint Class 3

Value	Device Type
4	Network Connectivity
5–255	Reserved

LLDP-MED Network Policies TLV

A network policy in the context of LLDP-MED is a device's VLAN configuration and associated Layer 2 and Layer 3 configurations.

LLDP-MED network policies TLV include:

- VLAN ID
- VLAN tagged or untagged status
- Layer 2 priority
- DSCP value

An integer represents the application type (the Type integer shown in the following table), which indicates a device function for which a unique network policy is defined. An individual LLDP-MED network policy TLV is generated for each application type that you specify with the CLI (XXAdvertising TLVs).


 **NOTE:** As shown in the following table, signaling is a series of control packets that are exchanged between an endpoint device and a network connectivity device to establish and maintain a connection. These signal packets might require a different network policy than the media packets for which a connection is made. In this case, configure the signaling application.

Table 37. Network Policy Applications

Type	Application	Description
0	Reserved	—
1	Voice	Specify this application type for dedicated IP telephony handsets and other appliances supporting interactive voice services.
2	Voice Signaling	Specify this application type only if voice control packets use a separate network policy than voice data.
3	Guest Voice	Specify this application type for a separate limited voice service for guest users with their own IP telephony handsets and other appliances supporting interactive voice services.
4	Guest Voice Signaling	Specify this application type only if guest voice control packets use a separate network policy than voice data.
5	Softphone Voice	Specify this application type only if guest voice control packets use a separate network policy than voice data.

Type	Application	Description
6	Video Conferencing	Specify this application type for dedicated video conferencing and other similar appliances supporting real-time interactive video.
7	Streaming Video	Specify this application type for dedicated video conferencing and other similar appliances supporting real-time interactive video.
8	Video Signaling	Specify this application type only if video control packets use a separate network policy than video data.
9–255	Reserved	—




Figure 74. LLDP-MED Policies TLV

Extended Power via MDI TLV

The extended power via MDI TLV enables advanced PoE management between LLDP-MED endpoints and network connectivity devices.

Advertise the extended power via MDI on all ports that are connected to an 802.3af powered, LLDP-MED endpoint device.

- **Power Type** — there are two possible power types: power source entity (PSE) or power device (PD). The Dell Networking system is a PSE, which corresponds to a value of 0, based on the TIA-1057 specification.
- **Power Source** — there are two possible power sources: primary and backup. The Dell Networking system is a primary power source, which corresponds to a value of 1, based on the TIA-1057 specification.
- **Power Priority** — there are three possible priorities: Low, High, and Critical. On Dell Networking systems, the default power priority is **High**, which corresponds to a value of 2 based on the TIA-1057 specification. You can configure a different power priority through the CLI. Dell Networking also honors the power priority value the powered device sends; however, the CLI configuration takes precedence.
- **Power Value** — Dell Networking advertises the maximum amount of power that can be supplied on the port. By default the power is **15.4W**, which corresponds to a power value of 130, based on the TIA-1057 specification. You can advertise a different power value using the `max-milliwatts` option with the `power inline auto | static` command. Dell Networking also honors the power value (power requirement) the powered device sends when the port is configured for `power inline auto`.




Figure 75. Extended Power via MDI TLV

Configure LLDP

Configuring LLDP is a two-step process.

1. Enable LLDP globally.
2. Advertise TLVs out of an interface.

Related Configuration Tasks

- [Viewing the LLDP Configuration](#)
- [Viewing Information Advertised by Adjacent LLDP Agents](#)
- [Configuring LLDPDU Intervals](#)
- [Configuring Transmit and Receive Mode](#)
- [Configuring a Time to Live](#)
- [Debugging LLDP](#)

Important Points to Remember

- LLDP is enabled by default.
- Dell Networking systems support up to eight neighbors per interface.
- Dell Networking systems support a maximum of 8000 total neighbors per system. If the number of interfaces multiplied by eight exceeds the maximum, the system does not configure more than 8000.
- INTERFACE level configurations override all CONFIGURATION level configurations.
- LLDP is not hitless.

LLDP Compatibility

- Spanning tree and force10 ring protocol “blocked” ports allow LLDPDUs.

CONFIGURATION versus INTERFACE Configurations

All LLDP configuration commands are available in PROTOCOL LLDP mode, which is a sub-mode of the CONFIGURATION mode and INTERFACE mode.

- Configurations made at the CONFIGURATION level are global; that is, they affect all interfaces on the system.
- Configurations made at the INTERFACE level affect only the specific interface; they override CONFIGURATION level configurations.

Example of the protocol lldp Command (CONFIGURATION Level)

```
R1(conf)#protocol lldp
R1(conf-lldp)#?
advertise      Advertise TLVs
dcbx           Configure Dcbx Parameters
disable Disable LLDP protocol globally
end            Exit from configuration mode
exit           Exit from LLDP configuration mode
fcoe           Configure priority bits for FCoE traffic
hello          LLDP hello configuration
iscsi          Configure priority bits for ISCSI traffic
mode           LLDP mode configuration (default = rx and tx)
multiplier     LLDP multiplier configuration
no             Negate a command or set its defaults
show           Show LLDP configuration
R1(conf-lldp)#exit
R1(conf)#interface tengigabitethernet 1/31
R1(conf-if-te-1/31)#protocol lldp
R1(conf-if-te-1/31-lldp)#?
advertise      Advertise TLVs
dcbx           Configure Dcbx Parameters
disable        Disable LLDP protocol on this interface
end            Exit from configuration mode
exit           Exit from LLDP configuration mode
hello          LLDP hello configuration
mode           LLDP mode configuration (default = rx and tx)
multiplier     LLDP multiplier configuration
no             Negate a command or set its defaults
show           Show LLDP configuration
no             Negate a command or set its defaults
show           Show LLDP configuration
R1(conf-if-te-1/31-lldp)#
```

Enabling LLDP

LLDP is enabled by default. Enable and disable LLDP globally or per interface. If you enable LLDP globally, all UP interfaces send periodic LLDPDUs.

To enable LLDP, use the following command.

1. Enter Protocol LLDP mode.
CONFIGURATION or INTERFACE mode

```
protocol lldp
```

2. Enable LLDP.
PROTOCOL LLDP mode

```
no disable
```

Disabling and Undoing LLDP

To disable or undo LLDP, use the following command.

- Disable LLDP globally or for an interface.
disable

To undo an LLDP configuration, precede the relevant command with the keyword `no`.

Advertising TLVs

You can configure the system to advertise TLVs out of all interfaces or out of specific interfaces.

- If you configure the system globally, all interfaces send LLDPDUs with the specified TLVs.
- If you configure an interface, only the interface sends LLDPDUs with the specified TLVs.
- If you configure LLDP both globally and at interface level, the interface level configuration overrides the global configuration.

To advertise TLVs, use the following commands.

1. Enter LLDP mode.

CONFIGURATION or INTERFACE mode

```
protocol lldp
```

2. Advertise one or more TLVs.

PROTOCOL LLDP mode

```
advertise {management-tlv | dot1-tlv | dot3-tlv | med | dcbx-appln-tlv |  
dcbx-tlv | interface-port-desc}
```

Include the keyword for each TLV you want to advertise.

- For management TLVs: `system-capabilities`, `system-description`.
- For 802.1 TLVs: `port-protocol-vlan-id`, `port-vlan-id` .
- For 802.3 TLVs: `max-frame-size`.
- For TIA-1057 TLVs:
 - `guest-voice`
 - `guest-voice-signaling`
 - `location-identification`
 - `power-via-mdi`
 - `softphone-voice`
 - `streaming-video`
 - `video-conferencing`
 - `video-signaling`
 - `voice`
 - `voice-signaling`

In the following example, LLDP is enabled globally. R1 and R2 are transmitting periodic LLDPDUs that contain management, 802.1, and 802.3 TLVs.




Figure 76. Configuring LLDP

Viewing the LLDP Configuration

To view the LLDP configuration, use the following command.

- Display the LLDP configuration.
CONFIGURATION or INTERFACE mode

```
show config
```

Example of Viewing LLDP Global Configurations

Example of Viewing LLDP Interface Configurations

```

R1(conf)#protocol lldp
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  hello 10
  no disable
R1(conf-lldp)#

```

```

R1(conf-lldp)#exit
R1(conf)#interface gigabitethernet 1/31
R1(conf-if-gi-1/31)#show config
!
interface GigabitEthernet 1/31
  no ip address
  switchport
  no shutdown
R1(conf-if-gi-1/31)#protocol lldp
R1(conf-if-gi-1/31-lldp)#show config
!
  protocol lldp
R1(conf-if-gi-1/31-lldp)#

```

Viewing Information Advertised by Adjacent LLDP Agents

To view brief information about adjacent devices or to view all the information that neighbors are advertising, use the following commands.

- Display brief information about adjacent devices.
`show lldp neighbors`
- Display all of the information that neighbors are advertising.
`show lldp neighbors detail`

Example of Viewing Brief Information Advertised by Neighbors

Example of Viewing Details Advertised by Neighbors

```
R1(conf-if-te-1/31-lldp)#end
R1(conf-if-te-1/31)#do show lldp neighbors
Loc PortID Rem Host Name Rem Port Id      Rem Chassis Id
-----
Te 0/2      -                00:00:c9:b1:3b:82  00:00:c9:b1:3b:82
Te 0/3      -                00:00:c9:ad:f6:12  00:00:c9:ad:f6:12
```

```
Dell#show lldp neighbors detail
=====
Local Interface Te 0/2 has 1 neighbor
  Total Frames Out: 16843
  Total Frames In: 17464
  Total Neighbor information Age outs: 0
  Total Multiple Neighbors Detected: 0
  Total Frames Discarded: 0
  Total In Error Frames: 0
  Total Unrecognized TLVs: 0
  Total TLVs Discarded: 0
  Next packet will be sent after 16 seconds
  The neighbors are given below:
-----

  Remote Chassis ID Subtype: Mac address (4)
  Remote Chassis ID: 00:00:c9:b1:3b:82
  Remote Port Subtype: Mac address (3)
  Remote Port ID: 00:00:c9:b1:3b:82
  Local Port ID: TenGigabitEthernet 0/2
  Locally assigned remote Neighbor Index: 7
  Remote TTL: 120
  Information valid for next 105 seconds
  Time since last information change of this neighbor: 1d21h56m
  Remote System Desc: Emulex OneConnect 10Gb Multi function Adapter
  Existing System Capabilities: Station only
  Enabled System Capabilities: Station only
-----

=====
Local Interface Te 0/3 has 1 neighbor
  Total Frames Out: 39165
  Total Frames In: 40650
  Total Neighbor information Age outs: 0
  Total Multiple Neighbors Detected: 0
  Total Frames Discarded: 0
```

```
Total In Error Frames: 0
Total Unrecognized TLVs: 0
Total TLVs Discarded: 0
Next packet will be sent after 4 seconds
The neighbors are given below:
```

```
Remote Chassis ID Subtype: Mac address (4)
Remote Chassis ID: 00:00:c9:ad:f6:12
Remote Port Subtype: Mac address (3)
Remote Port ID: 00:00:c9:ad:f6:12
Local Port ID: TenGigabitEthernet 0/3
```

Configuring LLDPDU Intervals

LLDPDUs are transmitted periodically; the default interval is **30 seconds**.

To configure LLDPDU intervals, use the following command.

- Configure a non-default transmit interval.
CONFIGURATION mode or INTERFACE mode

```
hello
```

Example of Viewing LLDPDU Intervals

```
R1(conf)#protocol lldp
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#mode ?
rx          Rx only
tx          Tx only
R1(conf-lldp)#mode tx
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  mode tx
  no disable
R1(conf-lldp)#no mode
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#
```

Configuring Transmit and Receive Mode

After you enable LLDP, Dell Networking systems transmit *and* receive LLDPDUs by default. To configure the system to transmit or receive only and return to the default, use the following commands.

- Transmit only.
CONFIGURATION mode or INTERFACE mode

```
mode tx
```

- Receive only.
CONFIGURATION mode or INTERFACE mode

```
mode rx
```

- Return to the default setting.
CONFIGURATION mode or INTERFACE mode

```
no mode
```

Example of Configuring a Single Mode

```
R1(conf)#protocol lldp
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#mode ?
rx          Rx only
tx          Tx only
R1(conf-lldp)#mode tx
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  mode tx
  no disable
R1(conf-lldp)#no mode
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#
```


Configuring a Time to Live

The information received from a neighbor expires after a specific amount of time (measured in seconds) called a time to live (TTL).

The TTL is the product of the LLDPDU transmit interval (hello) and an integer called a multiplier. The default multiplier is **4**, which results in a default TTL of 120 seconds.

- Adjust the TTL value.
CONFIGURATION mode or INTERFACE mode.

```
multiplier
```

- Return to the default multiplier value.
CONFIGURATION mode or INTERFACE mode.

```
no multiplier
```

Example of the `multiplier` Command to Configure Time to Live

```
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#multiplier ?
<2-10>                Multiplier (default=4)
R1(conf-lldp)#multiplier 5
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  multiplier 5
  no disable
R1(conf-lldp)#no multiplier
R1(conf-lldp)#show config
!
protocol lldp
  advertise dot1-tlv port-protocol-vlan-id port-vlan-id
  advertise dot3-tlv max-frame-size
  advertise management-tlv system-capabilities system-description
  no disable
R1(conf-lldp)#
```

Debugging LLDP

You can view the TLVs that your system is sending and receiving. To view the TLVs, use the following commands.

- View a readable version of the TLVs.
`debug lldp brief`

- View a readable version of the TLVs plus a hexadecimal version of the entire LLDPDU.

```
debug lldp detail
```

```
Dell# debug lldp interface tengigabitethernet 1/2 packet detail tx
Dell#1w1d19h : Transmit timer blew off for local interface Gi 1/2
1w1d19h : Forming LLDP pkt to send out of interface Gi 1/2
1w1d19h : TLV: Chassis ID, Len: 7, Subtype: Mac address (4), Value: 00:01:e8:0d:b6:d6
1w1d19h : TLV: Port ID, Len: 20, Subtype: Interface name (5), Value: TenGigabitEthernet 1/2
1w1d19h : TLV: TTL, Len: 2, Value: 120
1w1d19h : TLV: SYS_DESC, Len: 207, Value: Dell Networks Real Time Operating System Software, Dell
Operating System Version: 1.0, Dell Application Software Version: E_MAIN4.7.5.276, Copyright (c)1999-Build
Time: Fri Oct 26 12:22:22 PDT 2007
1w1d19h : TLV: SYSTEM_CAPAB, Len: 4, Value: Existing: Repeater Bridge Router, Enabled: Repeater Bridge Router
1w1d19h : TLV: ENDOFPDU, Len: 0
1w1d19h : Sending LLDP pkt out of Gi 1/2 of length 270
1w1d19h : Packet dump:
1w1d19h : 01 80 c2 00 00 0e 00 01 e8 0d b7 3b 81 00 00 00
1w1d19h : 88 cc 02 07 04 00 01 e8 0d b6 d6 04 14 05 47 69
1w1d19h : 67 61 62 69 74 45 74 68 65 72 6e 65 74 20 31 2f
1w1d19h : 32 06 02 00 78 0c cf 46 6f 72 63 65 31 30 20 4e
1w1d19h : 65 74 77 6f 72 6b 73 20 52 65 61 6c 20 54 69 6d
1w1d19h : 65 20 4f 70 65 72 61 74 69 6e 67 20 53 79 73 74
1w1d19h : 65 6d 20 53 6f 66 74 77 61 72 65 2e 20 46 6f 72
1w1d19h : 63 65 31 30 20 4f 70 65 72 61 74 69 6e 67 20 53
1w1d19h : 79 73 74 65 6d 20 56 65 72 73 69 6f 6e 3a 20 31
1w1d19h : 2e 30 2e 20 46 6f 72 63 65 31 30 20 41 70 70 6c
1w1d19h : 69 63 61 74 69 6f 6e 20 53 6f 66 74 77 61 72 65
1w1d19h : 20 56 65 72 73 69 6f 6e 3a 20 45 5f 4d 41 49 4e
1w1d19h : 34 2e 37 2e 35 2e 32 37 36 2e 20 43 6f 70 79 72
1w1d19h : 69 67 68 74 20 28 63 29 20 31 39 39 39 2d 42 75
1w1d19h : 69 6c 64 20 54 69 6d 65 3a 20 46 72 69 20 4f 63
1w1d19h : 74 20 32 36 20 31 32 3a 32 32 3a 32 32 20 50 44
1w1d19h : 54 20 32 30 30 37 0e 04 00 16 00 16 00 00
1w1d19h : LLDP frame sent out successfully of Gi 1/2
1w1d19h : Started Transmit timer for Loc interface Gi 1/2 for time 30 sec
```

Figure 77. The debug lldp detail Command — LLDPDU Packet Dissection

Relevant Management Objects

Dell Networkings OS supports all IEEE 802.1AB MIB objects.

The following tables list the objects associated with:

- received and transmitted TLVs
- the LLDP configuration on the local agent
- IEEE 802.1AB Organizationally Specific TLVs
- received and transmitted LLDP-MED TLVs

Table 38. LLDP Configuration MIB Objects

MIB Object Category	LLDP Variable	LLDP MIB Object	Description
LLDP Configuration	adminStatus	lldpPortConfigAdminStatus	Whether you enable the local LLDP agent for transmit, receive, or both.
	msgTxHold	lldpMessageTxHoldMultiplier	Multiplier value.
	msgTxInterval	lldpMessageTxInterval	Transmit Interval value.
	rxInfoTTL	lldpRxInfoTTL	Time to live for received TLVs.
	txInfoTTL	lldpTxInfoTTL	Time to live for transmitted TLVs.
Basic TLV Selection	mibBasicTLVsTxEnable	lldpPortConfigTLVsTxEnable	Indicates which management TLVs are enabled for system ports.
	mibMgmtAddrInstanceTxEnable	lldpManAddrPortsTxEnable	The management addresses defined for the system and the ports through which they are enabled for transmission.
LLDP Statistics	statsAgeoutsTotal	lldpStatsRxPortAgeoutsTotal	Total number of times that a neighbor's information is deleted on the local system due to an rxInfoTTL timer expiration.
	statsFramesDiscardedTotal	lldpStatsRxPortFramesDiscardedTotal	Total number of LLDP frames received then discarded.
	statsFramesInErrorsTotal	lldpStatsRxPortFramesErrors	Total number of LLDP frames received on a port with errors.
	statsFramesInTotal	lldpStatsRxPortFramesTotal	Total number of LLDP frames received through the port.
	statsFramesOutTotal	lldpStatsTxPortFramesTotal	Total number of LLDP frames transmitted through the port.
	statsTLVsDiscardedTotal	lldpStatsRxPortTLVsDiscardedTotal	Total number of TLVs received then discarded.

MIB Object Category	LLDP Variable	LLDP MIB Object	Description
	statsTLVsUnrecognizedTotal	lldpStatsRxPortTLVsUnrecognizedTotal	Total number of all TLVs the local agent does not recognize.

Table 39. LLDP System MIB Objects

TLV Type	TLV Name	TLV Variable	System	LLDP MIB Object
1	Chassis ID	chassis ID subtype	Local	lldpLocChassisIdSubtype
			Remote	lldpRemChassisIdSubtype
		chassid ID	Local	lldpLocChassisId
			Remote	lldpRemChassisId
2	Port ID	port subtype	Local	lldpLocPortIdSubtype
			Remote	lldpRemPortIdSubtype
		port ID	Local	lldpLocPortId
			Remote	lldpRemPortId
4	Port Description	port description	Local	lldpLocPortDesc
			Remote	lldpRemPortDesc
5	System Name	system name	Local	lldpLocSysName
			Remote	lldpRemSysName
6	System Description	system description	Local	lldpLocSysDesc
			Remote	lldpRemSysDesc
7	System Capabilities	system capabilities	Local	lldpLocSysCapSupported
			Remote	lldpRemSysCapSupported
8	Management Address	enabled capabilities	Local	lldpLocSysCapEnabled
			Remote	lldpRemSysCapEnabled
		management address length	Local	lldpLocManAddrLen
			Remote	lldpRemManAddrLen
		management address subtype	Local	lldpLocManAddrSubtype

TLV Type	TLV Name	TLV Variable	System	LLDP MIB Object
			Remote	lldpRemManAddrSubtype
		management address	Local	lldpLocManAddr
			Remote	lldpRemManAddr
		interface numbering subtype	Local	lldpLocManAddrIfSubtype
			Remote	lldpRemManAddrIfSubtype
		interface number	Local	lldpLocManAddrIfId
			Remote	lldpRemManAddrIfId
		OID	Local	lldpLocManAddrOID
			Remote	lldpRemManAddrOID

Table 40. LLDP 802.1 Organizationally specific TLV MIB Objects

TLV Type	TLV Name	TLV Variable	System	LLDP MIB Object
127	Port-VLAN ID	PVID	Local	lldpXdot1LocPortVlanId
			Remote	lldpXdot1RemPortVlanId
127	Port and Protocol VLAN ID	port and protocol VLAN supported	Local	lldpXdot1LocProtoVlanSupported
			Remote	lldpXdot1RemProtoVlanSupported
		port and protocol VLAN enabled	Local	lldpXdot1LocProtoVlanEnabled
			Remote	lldpXdot1RemProtoVlanEnabled
		PPVID	Local	lldpXdot1LocProtoVlanId
			Remote	lldpXdot1RemProtoVlanId
127	VLAN Name	VID	Local	lldpXdot1LocVlanId
			Remote	lldpXdot1RemVlanId
		VLAN name length	Local	lldpXdot1LocVlanName

TLV Type	TLV Name	TLV Variable	System	LLDP MIB Object
			Remote	lldpXdot1RemVlanName
		VLAN name	Local	lldpXdot1LocVlanName
			Remote	lldpXdot1RemVlanName

Table 41. LLDP-MED System MIB Objects

TLV Sub-Type	TLV Name	TLV Variable	System	LLDP-MED MIB Object
1	LLDP-MED Capabilities	LLDP-MED Capabilities	Local	lldpXMedPortCapSupported
				lldpXMedPortConfigTLVsTxEnable
			Remote	lldpXMedRemCapSupported
				lldpXMedRemConfigTLVsTxEnable
			Local	lldpXMedLocDeviceClass
			Remote	lldpXMedRemDeviceClass
2	Network Policy	Application Type	Local	lldpXMedLocMediaPolicyAppType
			Remote	lldpXMedRemMediaPolicyAppType
			Local	lldpXMedLocMediaPolicyUnknown
			Remote	lldpXMedLocMediaPolicyUnknown
		Tagged Flag	Local	lldpXMedLocMediaPolicyTagged
			Remote	lldpXMedLocMediaPolicyTagged
			Local	lldpXMedLocMediaPolicyVlanID
			Remote	lldpXMedRemMediaPolicyVlanID
		VLAN ID	Local	lldpXMedLocMediaPolicyVlanID
			Remote	lldpXMedRemMediaPolicyVlanID

TLV Sub-Type	TLV Name	TLV Variable	System	LLDP-MED MIB Object
3	Location Identifier	L2 Priority	Local	IldpXMedLocMediaPolicyPriority
			Remote	IldpXMedRemMediaPolicyPriority
		DSCP Value	Local	IldpXMedLocMediaPolicyDscp
			Remote	IldpXMedRemMediaPolicyDscp
		Location Data Format	Local	IldpXMedLocLocationSubtype
			Remote	IldpXMedRemLocationSubtype
		Location ID Data	Local	IldpXMedLocLocationInfo
			Remote	IldpXMedRemLocationInfo
4	Extended Power via MDI	Power Device Type	Local	IldpXMedLocXPoEDeviceType
			Remote	IldpXMedRemXPoEDeviceType
		Power Source	Local	IldpXMedLocXPoEPSPowerSource
				IldpXMedLocXPoEPDPowerSource
			Remote	IldpXMedRemXPoEPSPowerSource
				IldpXMedRemXPoEPDPowerSource
		Power Priority	Local	IldpXMedLocXPoEPDPowerPriority
			Remote	IldpXMedRemXPoEPSPowerPriority
				IldpXMedRemXPoEPDPowerPriority

TLV Sub-Type	TLV Name	TLV Variable	System	LLDP-MED MIB Object
		Power Value	Local	lldpXMedLocXPoEPS EPortPowerAv
				lldpXMedLocXPoEP DPowerReq
			Remote	lldpXMedRemXPoEP SEPowerAv
				lldpXMedRemXPoEP DPowerReq

Microsoft Network Load Balancing

Network Load Balancing (NLB) is a clustering functionality that is implemented by Microsoft on Windows 2000 Server and Windows Server 2003 operating systems. NLB uses a distributed methodology or pattern to equally split and balance the network traffic load across a set of servers that are part of the cluster or group. NLB combines the servers into a single multicast group and attempts to use the standard multicast IP or unicast IP addresses, and MAC addresses for the transmission of network traffic. At the same time, it also uses a single virtual IP address for all clients as the destination IP address, which enables servers to join the same multicast group in a way that is transparent to the clients (the clients do not notice the addition of new servers to the group). The clients use a cluster IP address to connect to the server. The NLB functionality enables flooding of traffic over the VLAN ports (for unicast mode) or a subset of ports in a VLAN (for multicast mode) to avoid overloading and effective performance of the servers for optimal processing of data packets.

NLB functions in two modes, namely unicast mode and multicast mode. The cluster IP address and the associated cluster MAC address are configured in the NLB application running on the Windows Server. In the unicast mode, when the server IP address is attempted to be resolved to the MAC address using the ARP application, the switch determines whether the ARP reply, obtained from the server, is of an NLB type. The switch then maps the IP address (cluster IP) with the MAC address (cluster MAC address). In multicast mode, the cluster IP address is mapped to a cluster multicast MAC address that is configured using a static ARP CLI configuration command. After the NLB entry is learned, the traffic is forwarded to all the servers in the VLAN corresponding to the cluster virtual IP address.

NLB Unicast Mode Scenario

Consider a sample topology in which four servers, namely S1 through S4, are configured as a cluster or a farm. This set of servers is connected to a Layer 3 switch, which in turn is connected to the end-clients. The servers contain a single IP address (IP-cluster address of 172.16.2.20) and a single unicast MAC address (MAC-Cluster address of 00-bf-ac-10-00-01) for load-balancing. Because multiple ports of a switch cannot learn a single MAC address, the servers are assigned with MAC addresses of MAC-s1 to MAC-s4 respectively on S1 through S4 in addition to the MAC cluster address. All the servers of the cluster belong to the VLAN named VLAN1.

In unicast NLB mode, the following sequence of events occurs:

- The switch sends an ARP request to resolve the IP address to the cluster MAC address.
- The ARP servers send an ARP response with the MAC cluster address in the ARP header and a MAC address of MAC-s1/s2/s3/s4 (for servers S1 through S4) in the Ethernet header.
- The switch associates the IP address with the MAC cluster address with the last ARP response it obtains. Assume that in this case, the last ARP reply is obtained from MAC-s4.(assuming that the ARP response with MAC-s4 is received as the last one). The interface associated with server, S4, is added to the ARP table.
- With NLB feature enabled, after learning the NLB ARP entry, all the subsequent traffic is flooded on all ports in VLAN1.

With NLB, the data frame is forwarded to all the servers for them to perform load-balancing.

NLB Multicast Mode Scenario

Consider a sample topology in which four servers, namely S1 through S4, are configured as a cluster or a farm. This set of servers is connected to a Layer 3 switch, which in turn is connected to the end-clients. They contain a single multicast MAC address (MAC-Cluster: 03-00-5E-11-11-11).

In the multicast NLB mode, a static ARP configuration command is configured to associate the cluster IP address with a multicast cluster MAC address.

With multicast NLB mode, the data is forwarded to all the servers based on the port specified using the Layer 2 multicast command, which is the `mac-address-table static <multicast_mac> multicast vlan <vlan_id> output-range <port1>, <port2>` command in CONFIGURATION mode.

Limitations With Enabling NLB on Switches

The following limitations apply to switches on which you configure NLB:

- The NLB unicast mode uses switch flooding to transmit all packets to all the servers that are part of the VLAN. When a large volume of traffic is processed, the clustering performance might be impacted in a small way. This limitation is applicable to switches that perform unicast flooding in the software.
- The `ip vlan-flooding` command applies globally across the system and for all VLANs. In cases where the NLB is applicable and the ARP replies contain a discrepancy in the Ethernet SHA and ARP header SHA frames, a flooding of packets over the relevant VLAN occurs.
- The maximum number of concurrent clusters that is supported is eight.

Benefits and Working of Microsoft Clustering

Microsoft clustering allows multiple servers using Microsoft Windows to be represented by one MAC address and IP address in order to provide transparent failover or balancing. Dell Networking OS does not recognize server clusters by default; it must be configured to do so. When an ARP request is sent to a server cluster, either the active server or all the servers send a reply, depending on the cluster configuration. If the active server sends a reply, the Dell switch learns the active server's MAC address. If all servers reply, the switch registers only the last received ARP reply, and the switch learns one server's actual MAC address; the virtual MAC address is never learned. Because the virtual MAC address is never learned, traffic is forwarded to only one server rather than the entire cluster, and failover and balancing are not preserved.

To preserve failover and balancing, the switch forwards the traffic destined for the server cluster to all member ports in the VLAN connected to the cluster. To ensure that this happens, you must configure the `ip vlan-flooding` command on the Dell switch at the time that the Microsoft cluster is configured. The server MAC address is given in the Ethernet frame header of the ARP reply, while the virtual MAC address representing the cluster is given in the payload. Then, all the traffic destined for the cluster is flooded out of all member ports. Since all the servers in the cluster receive traffic, failover and balancing are preserved.

Enable and Disable VLAN Flooding

- The older ARP entries are overwritten whenever newer NLB entries are learned.
- All ARP entries, learned after the feature is enabled, are deleted when the feature is disabled, and RP2 triggers an ARP resolution. The feature is disabled with the `no ip vlan-flooding` command.
- When a port is added to the VLAN, the port automatically receives traffic if the feature is enabled. Old ARP entries are not deleted or updated.
- When a member port is deleted, its ARP entries are also deleted from the CAM.
- Port channels in the VLAN also receive traffic.
- There is no impact on the configuration from saving the configuration.
- The feature, if enabled, is displayed in the `show running-config` command output that displays the `ip vlan-flooding` CLI configuration. Apart from it, there is no indication of the enabling of this capability.

Configuring a Switch for NLB

To enable a switch for unicast NLB mode of functioning, perform the following steps:

Enter the `ip vlan-flooding` command to specify that all Layer 3 unicast routed data traffic, going through a VLAN member port, needs to be flooded across all the member ports of that VLAN. There might be some ARP table entries that are resolved through ARP packets, which had the Ethernet MAC SA different from the MAC information inside the ARP packet. This unicast data traffic flooding occurs only for those packets that use these ARP entries.

CONFIGURATION mode

```
ip vlan-flooding
```

To enable a switch for multicast NLB mode of functioning, perform the following steps:

1. In the multicast mode of NLB, add a static ARP entry by entering the `arp ip-address multicast-mac-address` command in Global configuration mode to associate an IP address with a multicast MAC address in the switch. This setting causes the multicast MAC address to be mapped to the cluster IP address for the NLB mode of operation of the switch.

INTERFACE mode

```
arp ip-address multicast-mac-address interface
```

2. Associate specific MAC or hardware addresses to VLANs.

CONFIGURATION mode

```
mac-address-table static multicast-mac-address vlan vlan-id output-range  
interface
```

Multicast Source Discovery Protocol (MSDP)

Multicast source discovery protocol (MSDP) is supported on the MXL switch platform.

Protocol Overview

MSDP is a Layer 3 protocol that connects IPv4 protocol-independent multicast-sparse mode (PIM-SM) domains. A domain in the context of MSDP is a contiguous set of routers operating PIM within a common boundary defined by an exterior gateway protocol, such as border gateway protocol (BGP).

Each rendezvous point (RP) peers with every other RP via the transmission control protocol (TCP). Through this connection, peers advertise the sources in their domain.

1. When an RP in a PIM-SM domain receives a PIM register message from a source, it sends a source-active (SA) message to MSDP peers, as shown in the following illustration.
2. Each MSDP peer receives and forwards the message to its peers away from the originating RP.
3. When an MSDP peer receives an SA message, it determines if there are any group members within the domain interested in any of the advertised sources. If there are, the receiving RP sends a join message to the originating RP, creating a shortest path tree (SPT) to the source.

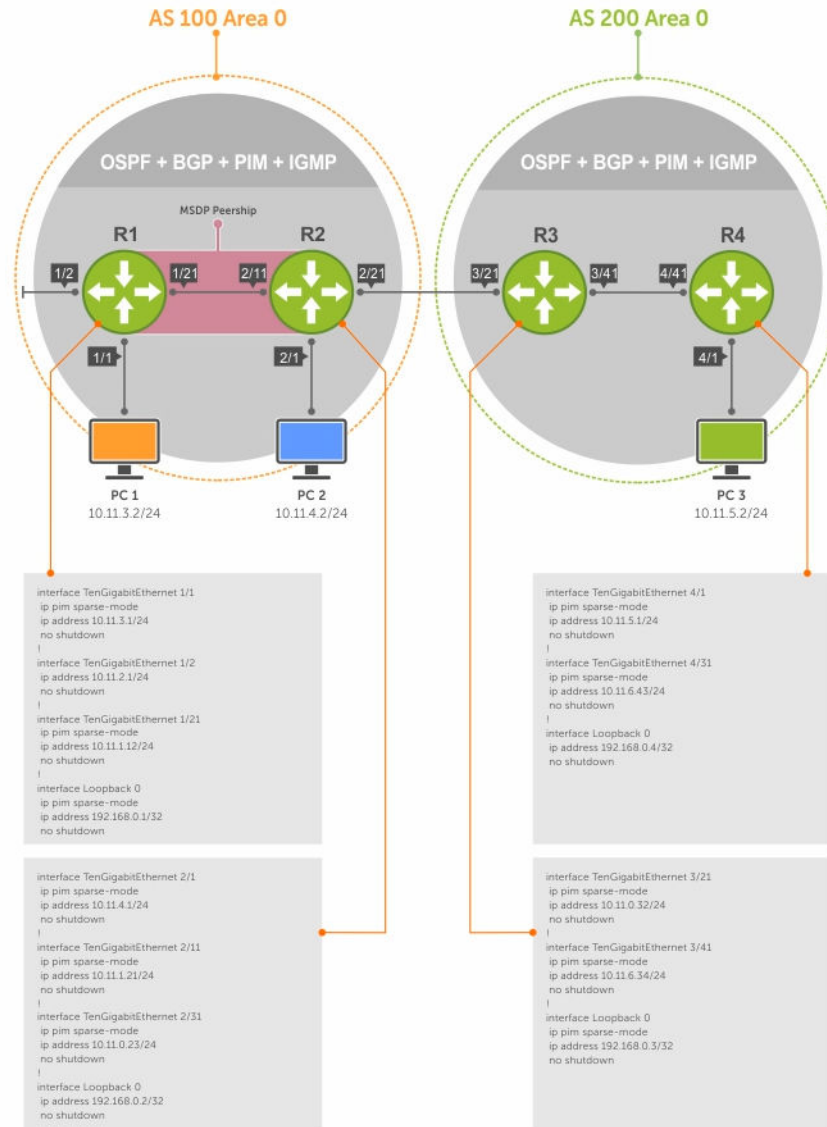


Figure 78. Multicast Source Discovery Protocol (MSDP)

RPs advertise each (S,G) in its domain in type, length, value (TLV) format. The total number of TLVs contained in the SA is indicated in the "Entry Count" field. SA messages are transmitted every 60 seconds, and immediately when a new source is detected.

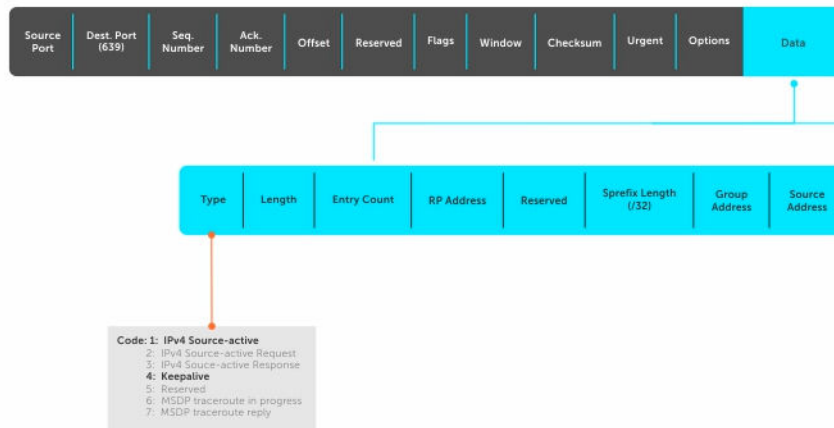


Figure 79. MSDP SA Message Format

Anycast RP

Using MSDP, anycast RP provides load sharing and redundancy in PIM-SM networks. Anycast RP allows two or more rendezvous points (RPs) to share the load for source registration and the ability to act as hot backup routers for each other.

Anycast RP allows you to configure two or more RPs with the same IP address on Loopback interfaces. The Anycast RP Loopback address are configured with a 32-bit mask, making it a host address. All downstream routers are configured to know that the Anycast RP Loopback address is the IP address of their local RP. IP routing automatically selects the closest RP for each source and receiver. Assuming that the sources are evenly spaced around the network, an equal number of sources register with each RP. Consequently, all the RPs in the network share the process of registering the sources equally. Because a source may register with one RP and receivers may join to a different RP, a method is needed for the RPs to exchange information about active sources. This information exchange is done with MSDP.

With Anycast RP, all the RPs are configured to be MSDP peers of each other. When a source registers with one RP, an SA message is sent to the other RPs informing them that there is an active source for a particular multicast group. The result is that each RP is aware of the active sources in the area of the other RPs. If any of the RPs fail, IP routing converges and one of the RPs becomes the active RP in more than one area. New sources register with the backup RP. Receivers join toward the new RP and connectivity is maintained.

Implementation Information

The Dell Networking operating system (OS) implementation of MSDP is in accordance with RFC 3618 and Anycast RP is in accordance with RFC 3446.

Configure the Multicast Source Discovery Protocol

Configuring MSDP is a four-step process.

1. Enable an exterior gateway protocol (EGP) with at least two routing domains. Refer to the following figures.

The [MSDP Sample Configurations](#) show the OSPF-BGP configuration used in this chapter for MSDP. Also, refer to [Open Shortest Path First \(OSPFv2 and OSPFv3\)](#) and [Border Gateway Protocol IPv4 \(BGPv4\)](#).

2. Configure PIM-SM within each EGP routing domain. Refer to the following figures.

The [MSDP Sample Configurations](#) show the PIM-SM configuration in this chapter for MSDP. Also, refer to [PIM Sparse-Mode \(PIM-SM\)](#).

3. [Enabling MSDP](#).
4. Peer the RPs in each routing domain with each other. Refer to [Enabling MSDP](#).

Related Configuration Tasks

The following lists related MSDP configuration tasks.

- [Enabling MSDP](#)
- [Manage the Source-Active Cache](#)
- [Accept Source-Active Messages that Fail the RFP Check](#)
- [Limiting the Source-Active Messages from a Peer](#)
- [Preventing MSDP from Caching a Local Source](#)
- [Preventing MSDP from Caching a Remote Source](#)
- [Preventing MSDP from Advertising a Local Source](#)
- [Terminating a Peership](#)
- [Clearing Peer Statistics](#)
- [Debugging MSDP](#)
- [Anycast RP](#)
- [MSDP Sample Configurations](#)

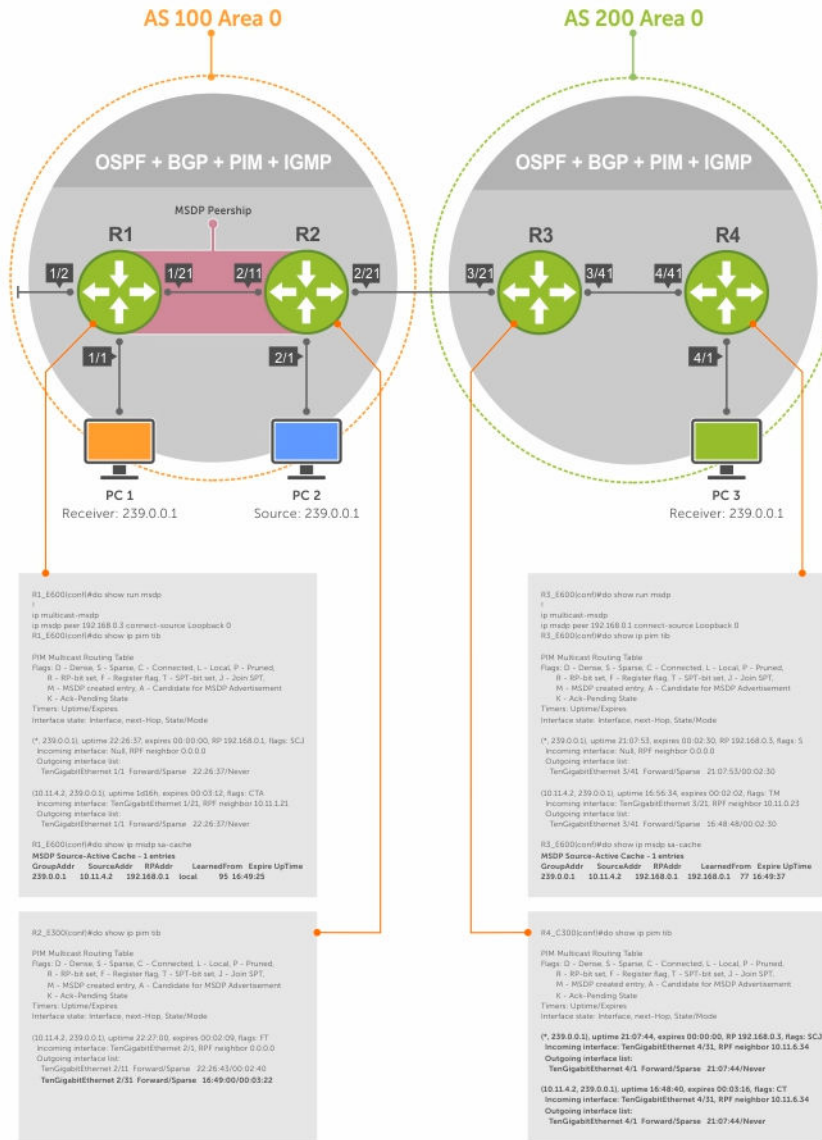


Figure 80. Configured Interfaces for MSDP

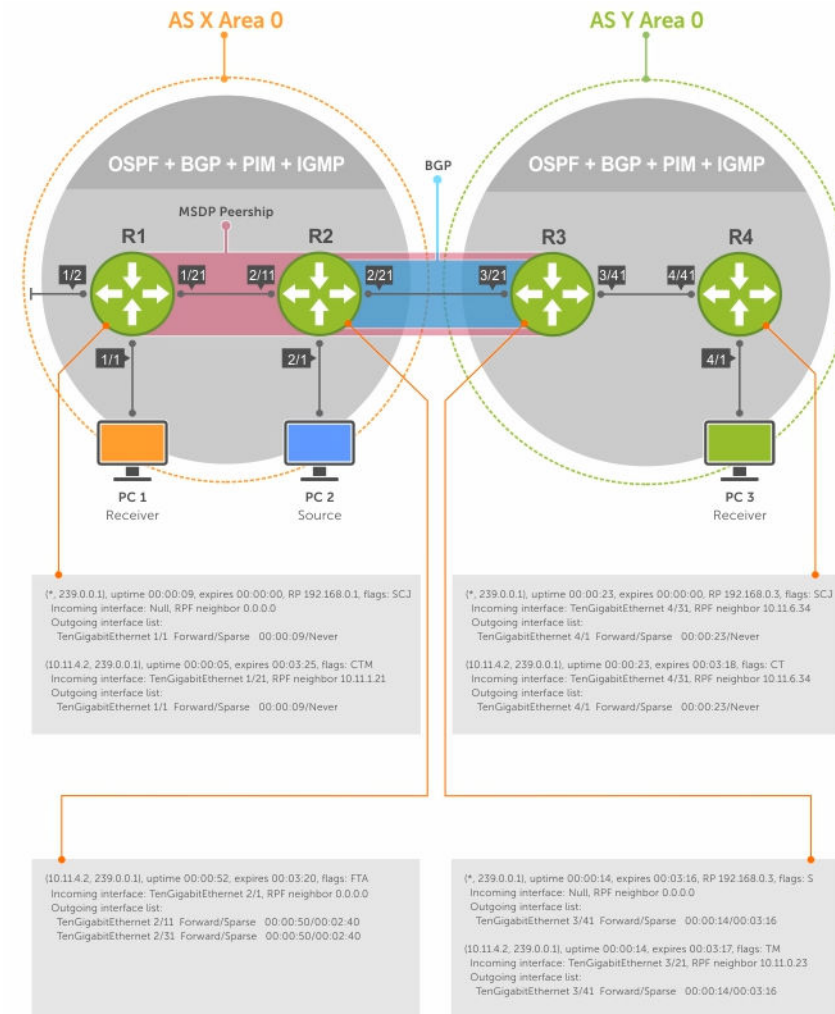


Figure 81. Configuring OSPF and BGP for MSDP

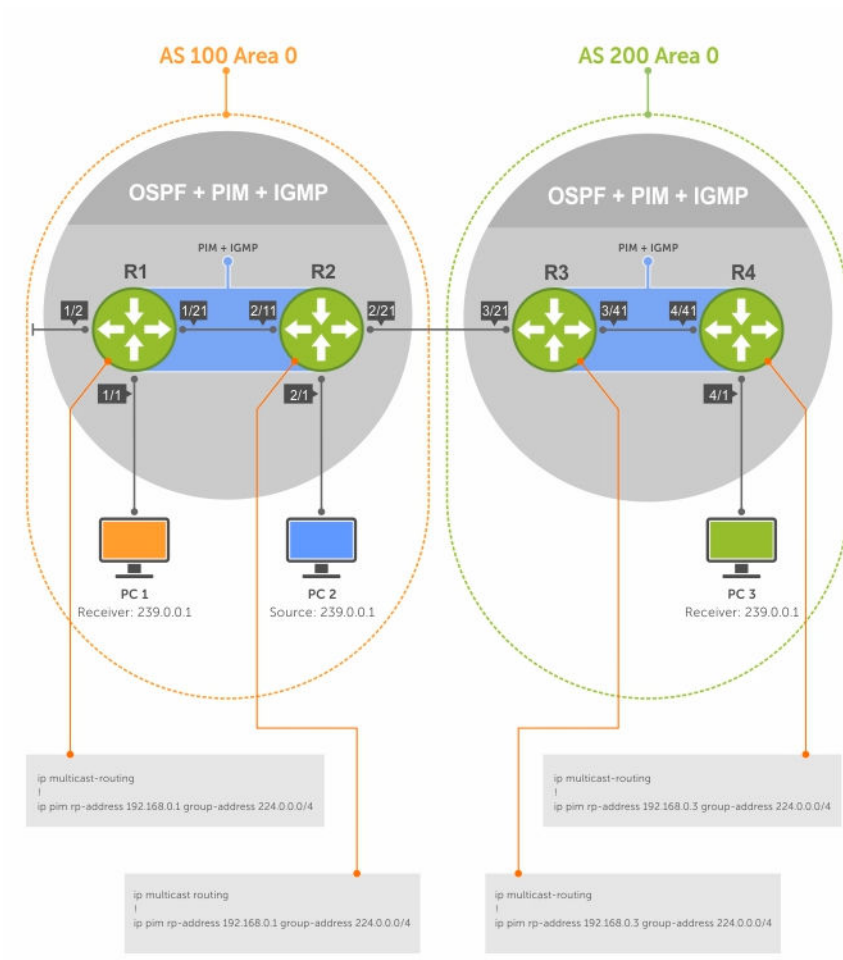


Figure 82. Configuring PIM in Multiple Routing Domains

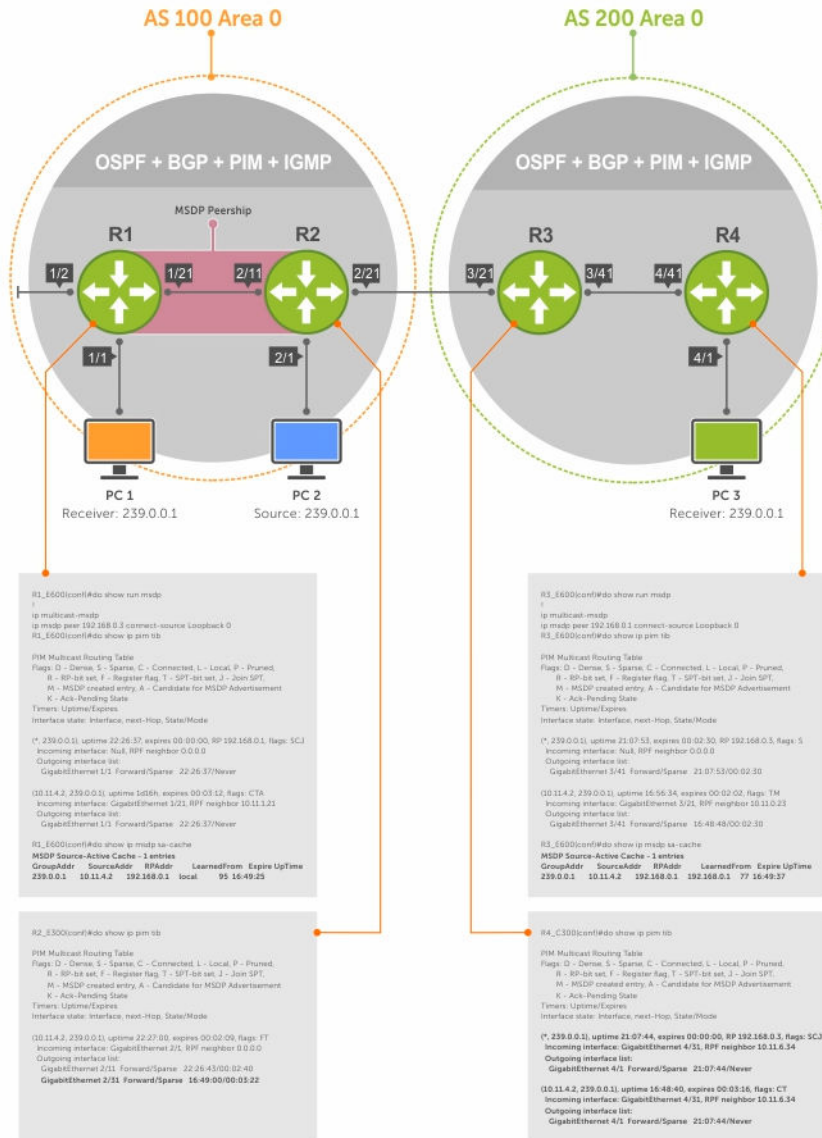


Figure 83. Configuring MSDP

Enabling MSDP

Enable MSDP by peering RPs in different administrative domains.

1. Enable MSDP.
CONFIGURATION mode
`ip multicast-msdp`
Peer PIM systems in different administrative domains.
CONFIGURATION mode

```
ip msdp peer connect-source
```

Example of Configuring MSDP

Example of Viewing Peer Information

```
R3_E600(conf)#ip multicast-msdp
R3_E600(conf)#ip msdp peer 192.168.0.1 connect-source Loopback 0
R3_E600(conf)#do show ip msdp summary

Peer Addr   Local Addr   State      Source   SA   Up/Down Description
192.168.0.1 192.168.0.3 Established Lo 0   1   00:05:29
```

To view details about a peer, use the `show ip msdp peer` command in EXEC privilege mode.

Multicast sources in remote domains are stored on the RP in the source-active cache (SA cache). The system does not create entries in the multicast routing table until there is a local receiver for the corresponding multicast group.

```
R3_E600#show ip msdp peer

Peer Addr: 192.168.0.1
Local Addr: 192.168.0.3(639) Connect Source: Lo 0
State: Established Up/Down Time: 00:15:20
Timers: KeepAlive 30 sec, Hold time 75 sec
SourceActive packet count (in/out): 8/0
SAs learned from this peer: 1
SA Filtering:
Input (S,G) filter: none
Output (S,G) filter: none
```

Manage the Source-Active Cache

Each SA-originating RP caches the sources inside its domain (domain-local), and the sources which it has learned from its peers (domain-remote).

By caching sources:

- domain-local receivers experience a lower join latency
- RPs can transmit SA messages periodically to prevent SA storms
- only sources that are in the cache are advertised in the SA to prevent transmitting multiple copies of the same source information

Viewing the Source-Active Cache

To view the source-active cache, use the following command.

- View the SA cache.
EXEC Privilege mode

```
show ip msdp sa-cache
```

Example of the show ip msdp sa-cache Command

```
R3_E600#show ip msdp sa-cache
MSDP Source-Active Cache - 1 entries
GroupAddr  SourceAddr  RPAAddr      LearnedFrom  Expire  UpTime
239.0.0.1  10.11.4.2   192.168.0.1  192.168.0.1  76     00:10:44
```

Limiting the Source-Active Cache

Set the upper limit of the number of active sources that the Dell Networking operating system caches. The default active source limit is 500K messages. When the total number of active sources reaches the specified limit, subsequent active sources are dropped even if they pass the reverse path forwarding (RPF) and policy check.

To limit the number of sources that SA cache stores, use the following command.

- Limit the number of sources that can be stored in the SA cache.
EXEC Privilege mode

```
show ip msdp sa-limit
```

If the total number of active sources is already larger than the limit when limiting is applied, the sources that are already in the system are not discarded. To enforce the limit in such a situation, use the `clear ip msdp sa-cache` command to clear all existing entries.

Clearing the Source-Active Cache

To clear the source-active cache, use the following command.

- Clear the SA cache of all, local, or rejected entries, or entries for a specific group.
CONFIGURATION mode

```
clear ip msdp sa-cache [group-address | local | rejected-sa]
```

Enabling the Rejected Source-Active Cache

To cache rejected sources, use the following command.

Active sources can be rejected because the RPF check failed, the SA limit is reached, the peer RP is unreachable, or the SA message has a format error.

- Cache rejected sources.
CONFIGURATION mode

```
ip msdp cache-rejected-sa
```

Accept Source-Active Messages that Fail the RFP Check

A default peer is a peer from which active sources are accepted even though they fail the RFP check.

Referring to the following illustrations:

- In Scenario 1, all MSPD peers are up.
- In Scenario 2, the peership between RP1 and RP2 is down, but the link (and routing protocols) between them is still up. In this case, RP1 learns all active sources from RP3, but the sources from RP2 and RP4 are rejected because the reverse path to these routers is through Interface A.
- In Scenario 3, RP3 is configured as a default MSPD peer for RP1 and so the RPF check is disregarded for RP3.
- In Scenario 4, RP1 has a default peer plus an access list. The list permits RP4 so the RPF check is disregarded for active sources from it, but RP5 (and all others because of the implicit deny all) are subject to the RPF check and fail, so those active sources are rejected.

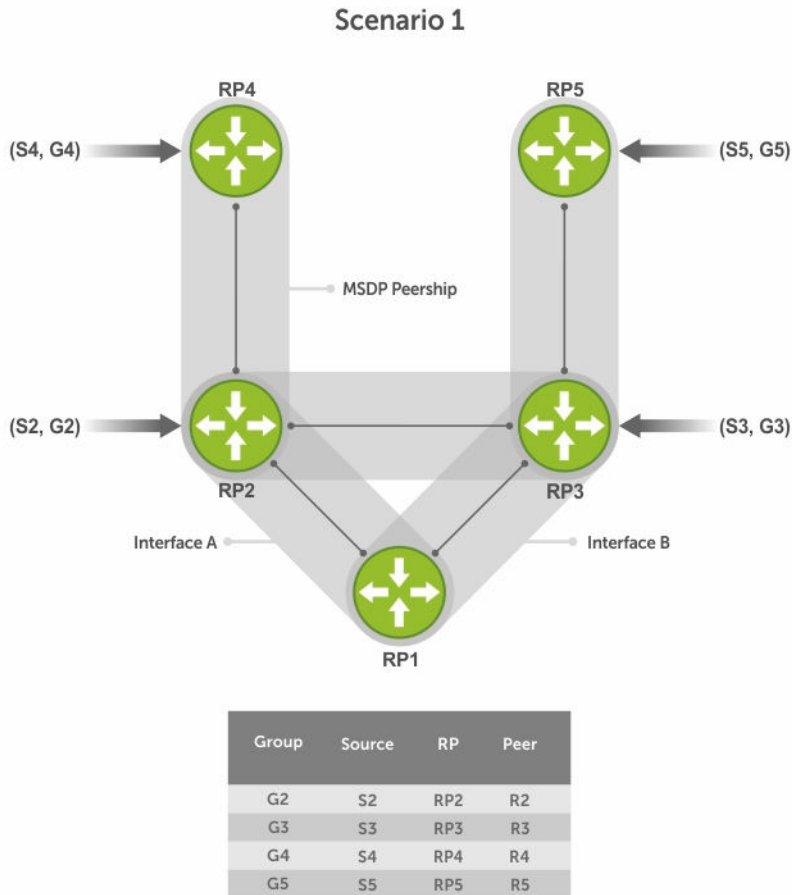



Figure 84. MSDP Default Peer, Scenario 1

Scenario 2



ip msdp default-peer Router 3

Group	Source	RP	Peer
G2	S2	RP2	R3 RPF-Fail
G3	S3	RP3	R3
G4	S4	RP4	R3 RPF-Fail
G5	S5	RP5	R3

Figure 85. MSDP Default Peer, Scenario 2

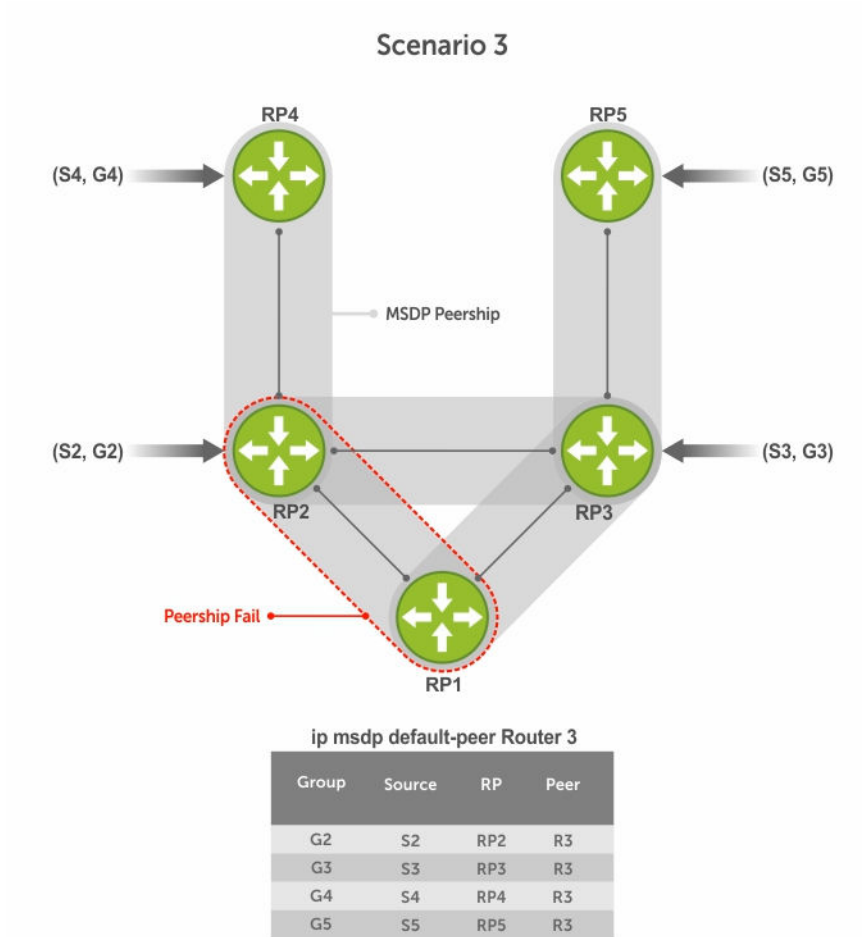
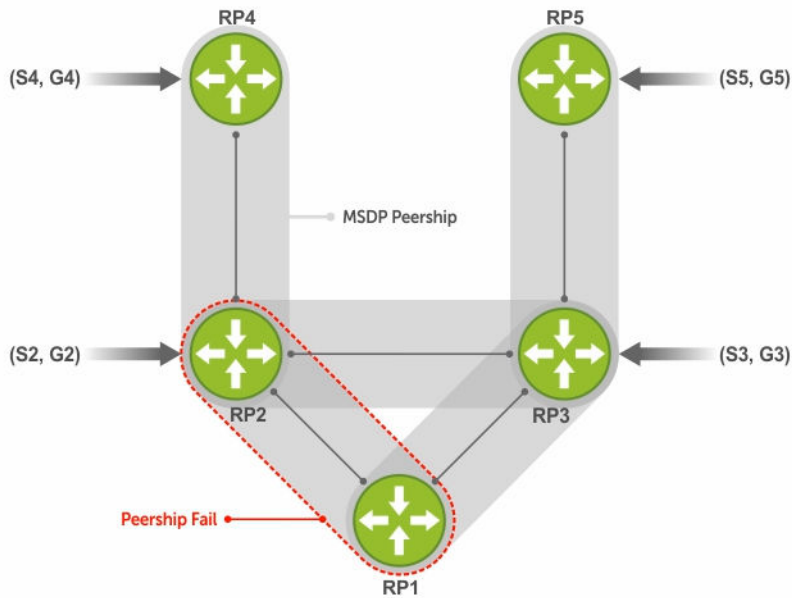


Figure 86. MSDP Default Peer, Scenario 3

Scenario 4



```
ip msdp default-peer Router 3 access-list list123
ip access-list list123
 permit RP4
 Deny RP5
```

Group	Source	RP	Peer
G2	S2	RP2	R3
G3	S3	RP3	R3 RPF-Fail
G4	S4	RP4	R3
G5	S5	RP5	R3

Figure 87. MSDP Default Peer, Scenario 4

Specifying Source-Active Messages

To specify messages, use the following command.

- Specify the forwarding-peer and originating-RP from which all active sources are accepted without regard for the RPF check.

CONFIGURATION mode

```
ip msdp default-peer ip-address list
```

If you do not specify an access list, the peer accepts all sources that peer advertises. All sources from RPs that the ACL denies are subject to the normal RPF check.

Example of the `ip msdp default-peer` Command and Viewing Denied Sources

```
Dell(conf)#ip msdp peer 10.0.50.2 connect-source Vlan 50
Dell(conf)#ip msdp default-peer 10.0.50.2 list fifty
```

```
Dell(conf)#ip access-list standard fifty
Dell(conf)#seq 5 permit host 200.0.0.50
```

```
Dell#ip msdp sa-cache
MSDP Source-Active Cache - 3 entries
GroupAddr  SourceAddr  RPAAddr  LearnedFrom  Expire  UpTime
229.0.50.2  24.0.50.2  200.0.0.50  10.0.50.2  73     00:13:49
229.0.50.3  24.0.50.3  200.0.0.50  10.0.50.2  73     00:13:49
229.0.50.4  24.0.50.4  200.0.0.50  10.0.50.2  73     00:13:49
```

```
Dell#ip msdp sa-cache rejected-sa
MSDP Rejected SA Cache
  3 rejected SAs received, cache-size 32766
UpTime  GroupAddr  SourceAddr  RPAAddr  LearnedFrom  Reason
00:33:18  229.0.50.64  24.0.50.64  200.0.1.50  10.0.50.2  Rpf-Fail
00:33:18  229.0.50.65  24.0.50.65  200.0.1.50  10.0.50.2  Rpf-Fail
00:33:18  229.0.50.66  24.0.50.66  200.0.1.50  10.0.50.2  Rpf-Fail
```

Limiting the Source-Active Messages from a Peer

To limit the source-active messages from a peer, use the following commands.

1. OPTIONAL: Store sources that are received after the limit is reached in the rejected SA cache.

CONFIGURATION mode

```
ip msdp cache-rejected-sa
```

2. Set the upper limit for the number of sources allowed from an MSDP peer.

CONFIGURATION mode

```
ip msdp peer peer-address sa-limit
```

The default limit is **100K**.

If the total number of sources received from the peer is already larger than the limit when this configuration is applied, those sources are not discarded. To enforce the limit in such a situation, first clear the SA cache.

Preventing MSDP from Caching a Local Source

You can prevent MSDP from caching an active source based on source and/or group. Because the source is not cached, it is not advertised to remote RPs.

1. OPTIONAL: Cache sources that are denied by the redistribute list in the rejected SA cache.

CONFIGURATION mode

```
ip msdp cache-rejected-sa
```

2. Prevent the system from caching local SA entries based on source and group using an extended ACL.

CONFIGURATION mode

```
ip msdp redistribute list
```

Example of Verifying the System is not Caching Local Sources

When you apply this filter, the SA cache is not affected immediately. When sources that are denied by the ACL time out, they are not refreshed. Until they time out, they continue to reside in the cache. To apply the redistribute filter to entries already present in the SA cache, first clear the SA cache. You may optionally store denied sources in the rejected SA cache.

```
R1_E600(conf)#do show run msdp
!
ip multicast-msdp
ip msdp peer 192.168.0.3 connect-source Loopback 0
ip msdp redistribute list mylocalfilter
ip msdp cache-rejected-sa 1000
R1_E600(conf)#do show run acl
!
ip access-list extended mylocalfilter
  seq 5 deny ip host 239.0.0.1 host 10.11.4.2
  seq 10 deny ip any any
R1_E600(conf)#do show ip msdp sa-cache
R1_E600(conf)#do show ip msdp sa-cache rejected-sa
MSDP Rejected SA Cache
  1 rejected SAs received, cache-size 1000
UpTime      GroupAddr  SourceAddr  RPAAddr      LearnedFrom  Reason
00:02:20   239.0.0.1  10.11.4.2   192.168.0.1  local        Redistribute
```

Preventing MSDP from Caching a Remote Source

To prevent MSDP from caching a remote source, use the following commands.

1. OPTIONAL: Cache sources that the SA filter denies in the rejected SA cache.
CONFIGURATION mode

```
ip msdp cache-rejected-sa
```

2. Prevent the system from caching remote sources learned from a specific peer based on source and group.

CONFIGURATION mode

```
ip msdp sa-filter list out peer list ext-acl
```

Example of Verifying the System is not Caching Remote Sources

As shown in the following example, R1 is advertising source 10.11.4.2. It is already in the SA cache of R3 when an ingress SA filter is applied to R3. The entry remains in the SA cache until it expires and is *not* stored in the rejected SA cache.

```
[Router 3]
R3_E600(conf)#do show run msdp
!
ip multicast-msdp
ip msdp peer 192.168.0.1 connect-source Loopback 0
ip msdp sa-filter in 192.168.0.1 list myremotefilter
R3_E600(conf)#do show run acl
```

```

!
ip access-list extended myremotefilter
 seq 5 deny ip host 239.0.0.1 host 10.11.4.2
R3_E600(conf)#do show ip msdp sa-cache
MSDP Source-Active Cache - 1 entries
GroupAddr SourceAddr RPAAddr LearnedFrom Expire UpTime
239.0.0.1 10.11.4.2 192.168.0.1 192.168.0.1 1 00:03:59
R3_E600(conf)#do show ip msdp sa-cache
R3_E600(conf)#
R3_E600(conf)#do show ip msdp peer

Peer Addr: 192.168.0.1
Local Addr: 0.0.0.0(639) Connect Source: Lo 0
State: Listening Up/Down Time: 00:01:19
Timers: KeepAlive 30 sec, Hold time 75 sec
SourceActive packet count (in/out): 0/0
SAs learned from this peer: 0
SA Filtering:
Input (S,G) filter: myremotefilter
Output (S,G) filter: none

```

Preventing MSDP from Advertising a Local Source

To prevent MSDP from advertising a local source, use the following command.

- Prevent an RP from advertising a source in the SA cache.
CONFIGURATION mode

```
ip msdp sa-filter list in peer list ext-acl
```

Example of Verifying the System is not Advertising Local Sources

In the following example, R1 stops advertising source 10.11.4.2. Because it is already in the SA cache of R3, the entry remains there until it expires.

```

[Router 1]
R1_E600(conf)#do show run msdp
!
ip multicast-msdp
ip msdp peer 192.168.0.3 connect-source Loopback 0
ip msdp sa-filter out 192.168.0.3 list mylocalfilter
R1_E600(conf)#do show run acl
!
ip access-list extended mylocalfilter
 seq 5 deny ip host 239.0.0.1 host 10.11.4.2
 seq 10 deny ip any any
R1_E600(conf)#do show ip msdp sa-cache
MSDP Source-Active Cache - 1 entries
GroupAddr SourceAddr RPAAddr LearnedFrom Expire UpTime
239.0.0.1 10.11.4.2 192.168.0.1 local 70 00:27:20
R3_E600(conf)#do show ip msdp sa-cache
MSDP Source-Active Cache - 1 entries
GroupAddr SourceAddr RPAAddr LearnedFrom Expire UpTime
239.0.0.1 10.11.4.2 192.168.0.1 192.168.0.1 1 00:10:29

[Router 3]
R3_E600(conf)#do show ip msdp sa-cache
R3_E600(conf)#

```

To display the configured SA filters for a peer, use the `show ip msdp peer` command from EXEC Privilege mode.

Logging Changes in Peership States

To log changes in peership states, use the following command.

- Log peership state changes.

CONFIGURATION mode

```
ip msdp log-adjacency-changes
```

Terminating a Peership

MSDP uses TCP as its transport protocol. In a peering relationship, the peer with the lower IP address initiates the TCP session, while the peer with the higher IP address listens on port 639.

- Terminate the TCP connection with a peer.

CONFIGURATION mode

```
ip msdp shutdown
```

Example of the Verifying that Peering State is Disabled

After the relationship is terminated, the peering state of the terminator is SHUTDOWN, while the peering state of the peer is INACTIVE.

```
[Router 3]
R3_E600(conf)#ip msdp shutdown 192.168.0.1
R3_E600(conf)#do show ip msdp peer

Peer Addr: 192.168.0.1
  Local Addr: 0.0.0.0(0) Connect Source: Lo 0
  State: Shutdown Up/Down Time: 00:00:18
  Timers: KeepAlive 30 sec, Hold time 75 sec
  SourceActive packet count (in/out): 0/0
  SAs learned from this peer: 0
  SA Filtering:
  Input (S,G) filter: myremotefilter
  Output (S,G) filter: none
[Router 1]
R1_E600(conf)#do show ip msdp peer

Peer Addr: 192.168.0.3
  Local Addr: 0.0.0.0(0) Connect Source: Lo 0
  State: Inactive Up/Down Time: 00:00:03
  Timers: KeepAlive 30 sec, Hold time 75 sec
  SourceActive packet count (in/out): 0/0
  SAs learned from this peer: 0
  SA Filtering:
```

Clearing Peer Statistics

To clear the peer statistics, use the following command.

- Reset the TCP connection to the peer and clear all peer statistics.
CONFIGURATION mode

```
clear ip msdp peer peer-address
```

Example of the `clear ip msdp peer` Command and Verifying Statistics are Cleared

```
R3_E600(conf)#do show ip msdp peer
```

```
Peer Addr: 192.168.0.1
  Local Addr: 192.168.0.3(639) Connect Source: Lo 0
  State: Established Up/Down Time: 00:04:26
  Timers: KeepAlive 30 sec, Hold time 75 sec
  SourceActive packet count (in/out): 5/0
  SAs learned from this peer: 0
  SA Filtering:
    Input (S,G) filter: myremotefilter
    Output (S,G) filter: none
```

```
R3_E600(conf)#do clear ip msdp peer 192.168.0.1
```

```
R3_E600(conf)#do show ip msdp peer
```

```
Peer Addr: 192.168.0.1
  Local Addr: 0.0.0.0(0) Connect Source: Lo 0
  State: Inactive Up/Down Time: 00:00:04
  Timers: KeepAlive 30 sec, Hold time 75 sec
  SourceActive packet count (in/out): 0/0
  SAs learned from this peer: 0
  SA Filtering:
    Input (S,G) filter: myremotefilter
    Output (S,G) filter: none
```

Debugging MSDP

To debug MSDP, use the following command.

- Display the information exchanged between peers.
CONFIGURATION mode

```
debug ip msdp
```

Example of the `debug ip msdp` Command

```
R1_E600(conf)#do debug ip msdp
```

```
All MSDP debugging has been turned on
```

```
R1_E600(conf)#03:16:08 : MSDP-0: Peer 192.168.0.3, sent Keepalive msg
03:16:09 : MSDP-0: Peer 192.168.0.3, rcvd Keepalive msg
03:16:27 : MSDP-0: Peer 192.168.0.3, sent Source Active msg
03:16:38 : MSDP-0: Peer 192.168.0.3, sent Keepalive msg
03:16:39 : MSDP-0: Peer 192.168.0.3, rcvd Keepalive msg
03:17:09 : MSDP-0: Peer 192.168.0.3, sent Keepalive msg
```

```
03:17:10 : MSDP-0: Peer 192.168.0.3, rcvd Keepalive msg
03:17:27 : MSDP-0: Peer 192.168.0.3, sent Source Active msg
Input (S,G) filter: none
Output (S,G) filter: none
```

MSDP with Anycast RP

Anycast RP uses MSDP with PIM-SM to allow more than one active group to use RP mapping.

PIM-SM allows only active groups to use RP mapping, which has several implications:

- **traffic concentration:** PIM-SM allows only one active group to RP mapping which means that all traffic for the group must, at least initially, travel over the same part of the network. You can load balance source registration between multiple RPs by strategically mapping groups to RPs, but this technique is less effective as traffic increases because preemptive load balancing requires prior knowledge of traffic distributions.
- **lack of scalable register decapsulation:** With only a single RP per group, all joins are sent to that RP regardless of the topological distance between the RP, sources, and receivers, and data is transmitted to the RP until the SPT switch threshold is reached.
- **slow convergence when an active RP fails:** When you configure multiple RPs, there can be considerable convergence delay involved in switching to the backup RP.

Anycast RP relieves these limitations by allowing multiple RPs per group, which can be distributed in a topologically significant manner according to the locations of the sources and receivers.

1. All the RPs serving a given group are configured with an identical anycast address.
2. Sources then register with the topologically closest RP.
3. RPs use MSDP to peer with each other using a unique address.

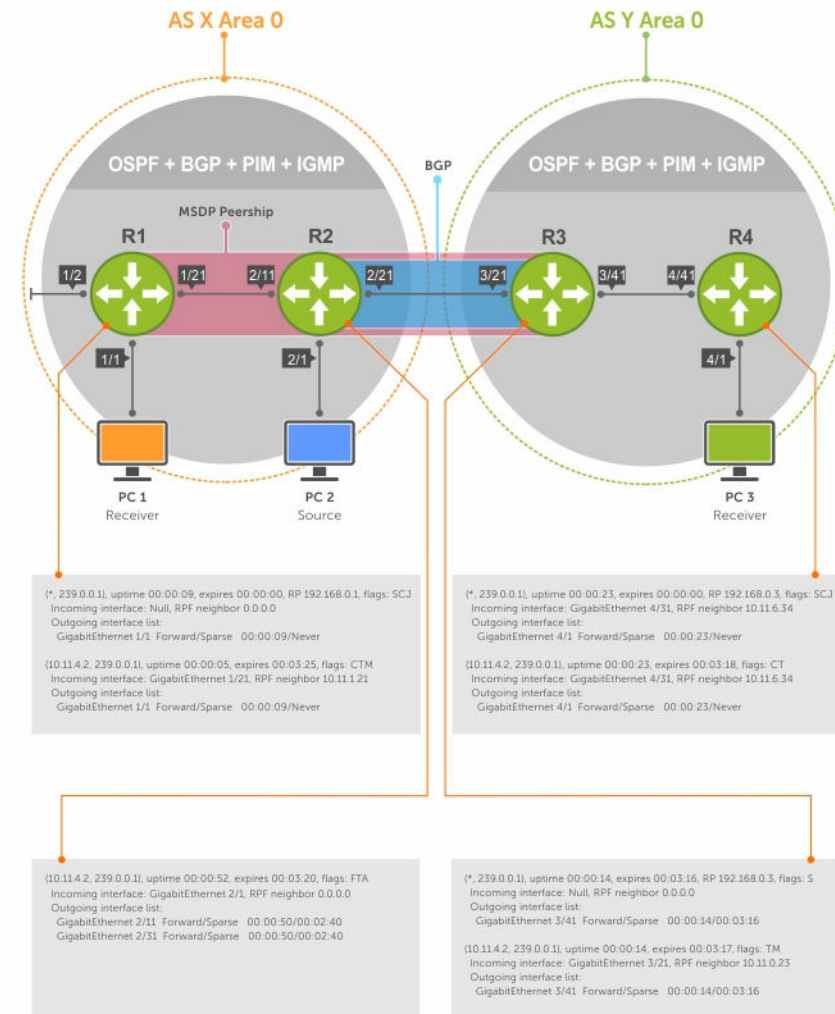


Figure 88. MSDP with Anycast RP

Configuring Anycast RP

To configure anycast RP, use the following commands.

1. In each routing domain that has multiple RPs serving a group, create a Loopback interface on each RP serving the group with the same IP address.

CONFIGURATION mode

```
interface loopback
```

2. Make this address the RP for the group.

CONFIGURATION mode

```
ip pim rp-address
```


3. In each routing domain that has multiple RPs serving a group, create another Loopback interface on each RP serving the group with a unique IP address.

CONFIGURATION mode

```
interface loopback
```

4. Peer each RP with every other RP using MSDP, specifying the unique Loopback address as the connect-source.

CONFIGURATION mode

```
ip msdp peer
```

5. Advertise the network of each of the unique Loopback addresses throughout the network.

ROUTER OSPF mode

```
network
```

Reducing Source-Active Message Flooding

RPs flood source-active messages to all of their peers away from the RP.

When multiple RPs exist within a domain, the RPs forward received active source information back to the originating RP, which violates the RFP rule. You can prevent this unnecessary flooding by creating a mesh-group. A mesh in this context is a topology in which each RP in a set of RPs has a peership with all other RPs in the set. When an RP is a member of the mesh group, it forwards active source information only to its peers outside of the group.

To create a mesh group, use the following command.

- Create a mesh group.

CONFIGURATION mode

```
ip msdp mesh-group
```

Specifying the RP Address Used in SA Messages

The default originator-id is the address of the RP that created the message. In the case of Anycast RP, there are multiple RPs all with the same address.

To use the (unique) address of another interface as the originator-id, use the following command.

- Use the address of another interface as the originator-id instead of the RP address.

CONFIGURATION mode

```
ip msdp originator-id
```

Example of R1 Configuration for MSDP with Anycast RP

Example of R2 Configuration for MSDP with Anycast RP

Example of R3 Configuration for MSDP with Anycast RP

```
ip multicast-routing
!
interface GigabitEthernet 1/1
 ip pim sparse-mode
 ip address 10.11.3.1/24
 no shutdown
```

```

!
interface GigabitEthernet 1/2
 ip address 10.11.2.1/24
 no shutdown
!
interface GigabitEthernet 1/21
 ip pim sparse-mode
 ip address 10.11.1.12/24
 no shutdown
!
interface Loopback 0
 ip pim sparse-mode
 ip address 192.168.0.1/32
 no shutdown
!
interface Loopback 1
 ip address 192.168.0.11/32
 no shutdown
!
router ospf 1
 network 10.11.2.0/24 area 0
 network 10.11.1.0/24 area 0
 network 10.11.3.0/24 area 0
 network 192.168.0.11/32 area 0
!
ip multicast-msdp
ip msdp peer 192.168.0.3 connect-source Loopback 1
ip msdp peer 192.168.0.22 connect-source Loopback 1
ip msdp mesh-group AS100 192.168.0.22
ip msdp originator-id Loopback 1!

ip pim rp-address 192.168.0.1 group-address 224.0.0.0/4

ip multicast-routing
!
interface GigabitEthernet 2/1
 ip pim sparse-mode
 ip address 10.11.4.1/24
 no shutdown
!
interface GigabitEthernet 2/11
 ip pim sparse-mode
 ip address 10.11.1.21/24
 no shutdown
!
interface GigabitEthernet 2/31
 ip pim sparse-mode
 ip address 10.11.0.23/24
 no shutdown
!
interface Loopback 0
 ip pim sparse-mode
 ip address 192.168.0.1/32
 no shutdown
!
interface Loopback 1
 ip address 192.168.0.22/32
 no shutdown
!
router ospf 1
 network 10.11.1.0/24 area 0
 network 10.11.4.0/24 area 0
 network 192.168.0.22/32 area 0

```

```

    redistribute static
    redistribute connected
    redistribute bgp 100
!
router bgp 100
    redistribute ospf 1
    neighbor 192.168.0.3 remote-as 200
    neighbor 192.168.0.3 ebgp-multihop 255
    neighbor 192.168.0.3 no shutdown
!
ip multicast-msdp
ip msdp peer 192.168.0.3 connect-source Loopback 1
ip msdp peer 192.168.0.11 connect-source Loopback 1
ip msdp mesh-group AS100 192.168.0.11
ip msdp originator-id Loopback 1
!
ip route 192.168.0.3/32 10.11.0.32
!
ip pim rp-address 192.168.0.1 group-address 224.0.0.0/4

ip multicast-routing
!
interface GigabitEthernet 3/21
    ip pim sparse-mode
    ip address 10.11.0.32/24
    no shutdown

interface GigabitEthernet 3/41
    ip pim sparse-mode
    ip address 10.11.6.34/24
    no shutdown
!
interface Loopback 0
    ip pim sparse-mode
    ip address 192.168.0.3/32
    no shutdown
!
router ospf 1
    network 10.11.6.0/24 area 0
    network 192.168.0.3/32 area 0
    redistribute static
    redistribute connected
    redistribute bgp 200
!
router bgp 200
    redistribute ospf 1
    neighbor 192.168.0.22 remote-as 100
    neighbor 192.168.0.22 ebgp-multihop 255
    neighbor 192.168.0.22 update-source Loopback 0
    neighbor 192.168.0.22 no shutdown
!
ip multicast-msdp
ip msdp peer 192.168.0.11 connect-source Loopback 0
ip msdp peer 192.168.0.22 connect-source Loopback 0
ip msdp sa-filter out 192.168.0.22
!
ip route 192.168.0.1/32 10.11.0.23
ip route 192.168.0.22/32 10.11.0.23
!
ip pim rp-address 192.168.0.3 group-address 224.0.0.0/4

```

MSDP Sample Configurations

The following examples show the running-configurations described in this chapter.

For more information, refer to the illustrations in the [Related Configuration Tasks](#) section.

MSDP Sample Configuration: R1 Running-Config

MSDP Sample Configuration: R2 Running-Config

MSDP Sample Configuration: R3 Running-Config

MSDP Sample Configuration: R4 Running-Config

```
ip multicast-routing
!
interface GigabitEthernet 1/1
 ip pim sparse-mode
 ip address 10.11.3.1/24
 no shutdown
!
interface GigabitEthernet 1/2
 ip address 10.11.2.1/24
 no shutdown
!
interface GigabitEthernet 1/21
 ip pim sparse-mode
 ip address 10.11.1.12/24
 no shutdown
!
interface Loopback 0
 ip pim sparse-mode
 ip address 192.168.0.1/32
 no shutdown
!
router ospf 1
 network 10.11.2.0/24 area 0
 network 10.11.1.0/24 area 0
 network 192.168.0.1/32 area 0
 network 10.11.3.0/24 area 0
!
ip multicast-msdp
ip msdp peer 192.168.0.3 connect-source Loopback 0
!
ip pim rp-address 192.168.0.1 group-address 224.0.0.0/4

ip multicast-routing
!
interface GigabitEthernet 2/1
 ip pim sparse-mode
 ip address 10.11.4.1/24
 no shutdown
!
interface GigabitEthernet 2/11
 ip pim sparse-mode
 ip address 10.11.1.21/24
 no shutdown
!
interface GigabitEthernet 2/31
```

```

    ip pim sparse-mode
    ip address 10.11.0.23/24
    no shutdown
!
interface Loopback 0
    ip address 192.168.0.2/32
    no shutdown
!
router ospf 1
    network 10.11.1.0/24 area 0
    network 10.11.4.0/24 area 0
    network 192.168.0.2/32 area 0
    redistribute static
    redistribute connected
    redistribute bgp 100
!
router bgp 100
    redistribute ospf 1
    neighbor 192.168.0.3 remote-as 200
    neighbor 192.168.0.3 ebgp-multihop 255
    neighbor 192.168.0.3 update-source Loopback 0
    neighbor 192.168.0.3 no shutdown
!
ip route 192.168.0.3/32 10.11.0.32
!
ip pim rp-address 192.168.0.1 group-address 224.0.0.0/4

ip multicast-routing
!
interface GigabitEthernet 3/21
    ip pim sparse-mode
    ip address 10.11.0.32/24
    no shutdown
!
interface GigabitEthernet 3/41
    ip pim sparse-mode
    ip address 10.11.6.34/24
    no shutdown
!
interface ManagementEthernet 0/0
    ip address 10.11.80.3/24
    no shutdown
!
interface Loopback 0
    ip pim sparse-mode
    ip address 192.168.0.3/32
    no shutdown
!
router ospf 1
    network 10.11.6.0/24 area 0
    network 192.168.0.3/32 area 0
    redistribute static
    redistribute connected
    redistribute bgp 200
!
router bgp 200
    redistribute ospf 1
    neighbor 192.168.0.2 remote-as 100
    neighbor 192.168.0.2 ebgp-multihop 255
    neighbor 192.168.0.2 update-source Loopback 0
    neighbor 192.168.0.2 no shutdown
!
ip multicast-msdp

```

```
ip msdp peer 192.168.0.1 connect-source Loopback 0
!
ip route 192.168.0.2/32 10.11.0.23

ip multicast-routing
!
interface GigabitEthernet 4/1
 ip pim sparse-mode
 ip address 10.11.5.1/24
 no shutdown
!
interface GigabitEthernet 4/22
 ip address 10.10.42.1/24
 no shutdown
!
interface GigabitEthernet 4/31
 ip pim sparse-mode
 ip address 10.11.6.43/24
 no shutdown
!
interface Loopback 0
 ip address 192.168.0.4/32
 no shutdown
!
router ospf 1
 network 10.11.5.0/24 area 0
 network 10.11.6.0/24 area 0
 network 192.168.0.4/32 area 0
!
ip pim rp-address 192.168.0.3 group-address 224.0.0.0/4
```

Multiple Spanning Tree Protocol (MSTP)

Multiple spanning tree protocol (MSTP) — specified in IEEE 802.1Q-2003 — is a rapid spanning tree protocol (RSTP)-based spanning tree variation that improves on per-VLAN spanning tree plus (PVST+).

Protocol Overview

MSTP allows multiple spanning tree instances and allows you to map many VLANs to one spanning tree instance to reduce the total number of required instances.

In contrast, PVST+ allows a spanning tree instance for each VLAN. This 1:1 approach is not suitable if you have many VLANs, because each spanning tree instance costs bandwidth and processing resources.

In the following illustration, three VLANs are mapped to two multiple spanning tree instances (MSTI). VLAN 100 traffic takes a different path than VLAN 200 and 300 traffic. The behavior demonstrates how you can use MSTP to achieve load balancing.




Figure 89. MSTP with Three VLANs Mapped to TWO Spanning Tree Instances

Spanning Tree Variations

The Dell Networking operating system (OS) supports four variations of spanning tree, as shown in the following table.

Table 42. Spanning Tree Variations

Dell Networking Term	IEEE Specification
Spanning Tree Protocol (STP)	802 .1d
Rapid Spanning Tree Protocol (RSTP)	802 .1w
Multicast Source Discovery Protocol (MSDP)	802 .1s
Per-VLAN Spanning Tree Plus (PVST+)	Third Party

Implementation Information

The following describes the MSTP implementation information.

- The Dell Networking OS MSTP implementation is based on IEEE 802.1Q-2003 and interoperates only with bridges that also use this standard implementation.
- MSTP is compatible with STP and RSTP.
- The Dell Networking OS supports only one MSTP region.
- When you enable MSTP, all ports in Layer 2 mode participate in MSTP.
- On the MXL switch, you can configure 64 MSTIs including the default instance 0 (CIST).

Configure Multiple Spanning Tree Protocol

Configuring multiple spanning tree is a four-step process.

1. Configure interfaces for Layer 2.
2. Place the interfaces in VLANs.
3. Enable the multiple spanning tree protocol.
4. Create multiple spanning tree instances and map VLANs to them.

Related Configuration Tasks

The following are the related configuration tasks for MSTP.

- [Creating Multiple Spanning Tree Instances](#)
- [Influencing MSTP Root Selection](#)
- [Interoperate with Non-Dell Networking OS Bridges](#)
- [Enable BPDU Filtering Globally](#)
- [Modifying Global Parameters](#)

- [Configuring an EdgePort](#)
- [Flush MAC Addresses after a Topology Change](#)
- [Debugging and Verifying MSTP Configurations](#)
- [Prevent Network Disruptions with BPDU Guard](#)
- [SNMP Traps for Root Elections and Topology Changes](#)

Enable Multiple Spanning Tree Globally

MSTP is not enabled by default. To enable MSTP globally, use the following commands. When you enable MSTP, all physical, VLAN, and port-channel interfaces that are enabled and in Layer 2 mode are automatically part of the MSTI 0.

- Within an MSTI, only one path from any bridge to any other bridge is enabled.
- Bridges block a redundant path by disabling one of the link ports.

1. Enter PROTOCOL MSTP mode.
CONFIGURATION mode

`protocol spanning-tree mstp`
2. Enable MSTP.
PROTOCOL MSTP mode

`no disable`

Example of Verifying MSTP is Enabled

To verify that MSTP is enabled, use the `show config` command in PROTOCOL MSTP mode.

```
Dell(conf)#protocol spanning-tree mstp
Dell(config-mstp)#show config
!
protocol spanning-tree mstp
no disable
Dell#
```

Creating Multiple Spanning Tree Instances

To create multiple spanning tree instances, use the following command. A single MSTI provides no more benefit than RSTP. To take full advantage of MSTP, create multiple MSTIs and map VLANs to them.

- Create an MSTI.
PROTOCOL MSTP mode

`msti`

Specify the keyword `vlan` then the VLANs that you want to participate in the MSTI.

Example of the `msti` Command

Example of Viewing MSTP Port States

```
Dell(conf)#protocol spanning-tree mstp
Dell(conf-mstp)#msti 1 vlan 100
Dell(conf-mstp)#msti 2 vlan 200-300
Dell(conf-mstp)#show config
!
protocol spanning-tree mstp
  no disable
  MSTI 1 VLAN 100
  MSTI 2 VLAN 200-300
```

All bridges in the MSTP region must have the same VLAN-to-instance mapping.

To view which instance a VLAN is mapped to, use the `show spanning-tree mst vlan` command from EXEC Privilege mode.

To view the forwarding/discarding state of the ports participating in an MSTI, use the `show spanning-tree msti` command from EXEC Privilege mode.

```
Dell#show spanning-tree msti 1
MSTI 1 VLANs mapped 100
Root Identifier has priority 32768, Address 0001.e806.953e
Root Bridge hello time 2, max age 20, forward delay 15, max hops 19
Bridge Identifier has priority 32768, Address 0001.e80d.b6d6
Configured hello time 2, max age 20, forward delay 15, max hops 20
Current root has priority 32768, Address 0001.e806.953e
Number of topology changes 2, last change occurred 1d2h ago on TenGig 1/21
Port 374 (TenGigabitEthernet 1/21) is root Forwarding
Port path cost 2000, Port priority 128, Port Identifier 128.374
Designated root has priority 32768, address 0001.e806.953e
Designated bridge has priority 32768, address 0001.e806.953e
Designated port id is 128.374, designated path cost 2000
Number of transitions to forwarding state 1
BPDU (MRecords): sent 93671, received 46843
The port is not in the Edge port mode, bpdu filter is disabled
Port 384 (TenGigabitEthernet 1/31) is alternate Discarding
Port path cost 2000, Port priority 128, Port Identifier 128.384
Designated root has priority 32768, address 0001.e806.953e
Designated bridge has priority 32768, address 0001.e809.c24a
Designated port id is 128.384, designated path cost 2000
Number of transitions to forwarding state 1
BPDU (MRecords): sent 39291, received 7547
The port is not in the Edge port mode, bpdu filter is disabled
```

Influencing MSTP Root Selection

MSTP determines the root bridge, but you can assign one bridge a lower priority to increase the probability that it becomes the root bridge.

To change the bridge priority, use the following command.

- Assign a number as the bridge priority.
PROTOCOL MSTP mode

`msti instance bridge-priority priority`

A lower number increases the probability that the bridge becomes the root bridge.

The range is from 0 to 61440, in increments of 4096.

The default is **32768**.

Example of Assigning and Verifying the Root Bridge Priority

By default, the simple configuration shown previously yields the same forwarding path for both MSTIs. The following example shows how R3 is assigned bridge priority 0 for MSTI 2, which elects a different root bridge than MSTI 2.

To view the bridge priority, use the `show config` command from PROTOCOL MSTP mode.

```
Dell(conf-mstp)#msti 2 bridge-priority 0

Dell(conf-mstp)#show config
!
protocol spanning-tree mstp
  MSTI 2 bridge-priority 0
Dell(conf-mstp)#
```


Interoperate with Non-Dell Networking OS Bridges

The Dell Networking OS supports only one MSTP region.

A region is a combination of three unique qualities:

- **Name** is a mnemonic string you assign to the region. The default region name is **null**.
- **Revision** is a 2-byte number. The default revision number is **0**.
- VLAN-to-instance mapping is the placement of a VLAN in an MSTI.

For a bridge to be in the same MSTP region as another, all three of these qualities must match exactly. The default values for name and revision match on all Dell Networking OS equipment. If you have non-Dell Networking OS equipment that participates in MSTP, ensure these values match on all the equipment.

 **NOTE:** Some non-Dell Networking OS equipment may implement a non-null default region name. SFTOS, for example, uses the Bridge ID, while others may use a MAC address.

Changing the Region Name or Revision

To change the region name or revision, use the following commands.

- Change the region name.
PROTOCOL MSTP mode

`name name`
- Change the region revision number.
PROTOCOL MSTP mode

`revision number`

The range is from 0 to 65535.

The default is **0**.

Example of the `name` Command


To view the current region name and revision, use the `show spanning-tree mst configuration` command from EXEC Privilege mode.

```
Dell(conf-mstp)#name my-mstp-region
Dell(conf-mstp)#exit
Dell(conf)#do show spanning-tree mst config
MST region name: my-mstp-region
Revision: 0
MSTI    VID
  1      100
  2      200-300
```

Modifying Global Parameters

The root bridge sets the values for forward-delay, hello-time, max-age, and max-hops and overwrites the values set on other MSTP bridges.

- **Forward-delay** — the amount of time an interface waits in the Listening state and the Learning state before it transitions to the Forwarding state.
- **Hello-time** — the time interval in which the bridge sends MSTP bridge protocol data units (BPDUs).
- **Max-age** — the length of time the bridge maintains configuration information before it refreshes that information by recomputing the MST topology.
- **Max-hops** — the maximum number of hops a BPDU can travel before a receiving switch discards it.

 **NOTE:** Dell Networking recommends that only experienced network administrators change MSTP parameters. Poorly planned modification of MSTP parameters can negatively affect network performance.

To change the MSTP parameters, use the following commands on the root bridge.

1. Change the forward-delay parameter.

```
PROTOCOL MSTP mode

forward-delay seconds
```


The range is from 4 to 30.

The default is **15 seconds**.

2. Change the hello-time parameter.

```
PROTOCOL MSTP mode

hello-time seconds
```

 **NOTE:** With large configurations (especially those configurations with more ports) Dell Networking recommends increasing the hello-time.

The range is from 1 to 10.

The default is **2 seconds**.

3. Change the max-age parameter.

PROTOCOL MSTP mode

```
max-age seconds
```

The range is from 6 to 40.

The default is **20 seconds**.

4. Change the max-hops parameter.

PROTOCOL MSTP mode

```
max-hops number
```

The range is from 1 to 40.

The default is **20**.

Example of the `forward-delay` Parameter

To view the current values for MSTP parameters, use the `show running-config spanning-tree mstp` command from EXEC privilege mode.

```
Dell(conf-mstp)#forward-delay 16
Dell(conf-mstp)#exit
Dell(conf)#do show running-config spanning-tree mstp
!
protocol spanning-tree mstp
  no disable
  name my-mstp-region
  MSTI 1 VLAN 100
  MSTI 2 VLAN 200-300
  forward-delay 16
  MSTI 2 bridge-priority 4096
Dell(conf)#
```

Enable BPDU Filtering Globally

The enabling of BPDU Filtering stops transmitting of BPDUs on the operational port fast enabled ports by default.

When BPDUs are received, the spanning tree is automatically prepared. By default global bpdu filtering is disabled.

Enable BPDU Filter globally to filter transmission of BPDU port fast enabled interfaces.

PROTOCOL MSTP mode

edge-port bpdu filter default

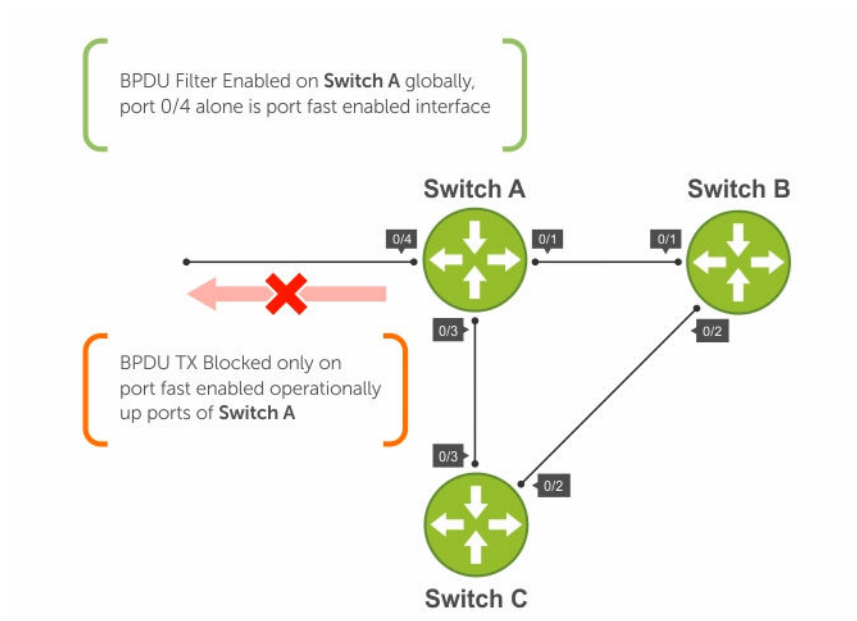


Figure 90. BPDUs Filtering Enabled Globally

Modifying the Interface Parameters

You can adjust two interface parameters to increase or decrease the probability that a port becomes a forwarding port.

- **Port cost** is a value that is based on the interface type. The greater the port cost, the less likely the port is selected to be a forwarding port.
- **Port priority** influences the likelihood that a port is selected to be a forwarding port in case that several ports have the same port cost.

The following lists the default values for port cost by interface.

Table 43. Default Values for Port Costs by Interface

Port Cost	Default Value
1000-Mb/s Ethernet interfaces	20000
40-Gigabit Ethernet interfaces	1400
10-Gigabit Ethernet interfaces	2000
Port Channel with one 10-Gigabit Ethernet interface	2000
Port Channel with one 40-Gigabit Ethernet interface	1400
Port Channel with two 10-Gigabit Ethernet interfaces	1800
Port Channel with two 40-Gigabit Ethernet interfaces	600

To change the port cost or priority of an interface, use the following commands.

1. Change the port cost of an interface.

INTERFACE mode

```
spanning-tree msti number cost cost
```

The range is from 0 to 200000.

For the default, refer to the default values shown in the table.

2. Change the port priority of an interface.

INTERFACE mode

```
spanning-tree msti number priority priority
```

The range is from 0 to 240, in increments of 16.

The default is **128**.

To view the current values for these interface parameters, use the `show config` command from INTERFACE mode.

Configuring an EdgePort

The EdgePort feature enables interfaces to begin forwarding traffic approximately 30 seconds sooner. In this mode, an interface forwards frames by default until it receives a BPDU that indicates that it should behave otherwise; it does not go through the Learning and Listening states. The `bpduguard shutdown-on-violation` option causes the interface hardware to be shut down when it receives a BPDU. When you implement only `bpduguard`, although the interface is placed in an Error Disabled state when receiving the BPDU, the physical interface remains up and spanning-tree drops packets in the hardware after a BPDU violation. BPDUs are dropped in the software after receiving the BPDU violation. This feature is the same as PortFast mode in spanning tree.



CAUTION: Configure EdgePort only on links connecting to an end station. EdgePort can cause loops if you enable it on an interface connected to a network.

To enable EdgePort on an interface, use the following command.

- Enable EdgePort on an interface.

INTERFACE mode

```
spanning-tree mstp edge-port [bpduguard | shutdown-on-violation]
```

Dell Networking OS Behavior: Regarding `bpduguard shutdown-on-violation` behavior:

- If the interface to be shut down is a port channel, all the member ports are disabled in the hardware.
- When you add a physical port to a port channel already in the Error Disable state, the new member port is also disabled in the hardware.
- When you remove a physical port from a port channel in the Error Disable state, the error disabled state is cleared on this physical port (the physical port is enabled in the hardware).
- You can clear the Error Disabled state with any of the following methods:

- * Use the `shutdown` command on the interface.
- * Disable the `shutdown-on-violation` command on the interface (using the `no spanning-tree mstp edge-port [bpduguard | [shutdown-on-violation]]` command).
- * Disable spanning tree on the interface (using the `no spanning-tree` command in INTERFACE mode).
- * Disabling global spanning tree (using the `no spanning-tree` command in CONFIGURATION mode).

Example of Enabling an EdgePort on an Interface

To verify that EdgePort is enabled, use the `show config` command from INTERFACE mode.

```
Dell(conf-if-gi-3/41)#spanning-tree mstp edge-port
Dell(conf-if-gi-3/41)#show config
!
interface GigabitEthernet 3/41
  no ip address
  switchport
  spanning-tree mstp edge-port
  spanning-tree MSTI 1 priority 144
  no shutdown
Dell(conf-if-gi-3/41)#
```

Flush MAC Addresses after a Topology Change

The Dell Networking OS has an optimized MAC address flush mechanism for RSTP, MSTP, and PVST+ that flushes addresses only when necessary, which allows for faster convergence during topology changes.

However, you may activate the flushing mechanism defined by 802.1Q-2003 using the `tc-flush-standard` command, which flushes MAC addresses after every topology change notification.

To view the enable status of this feature, use the `show running-config spanning-tree mstp` command from EXEC Privilege mode.

MSTP Sample Configurations

The running-configurations support the topology shown in the following illustration.

The configurations are from Dell Networking OS systems.




Figure 91. MSTP with Three VLANs Mapped to Two Spanning Tree Instances

Router 1 Running-Configuration

This example uses the following steps:

1. Enable MSTP globally and set the region name and revision map MSTP instances to the VLANs.
2. Assign Layer-2 interfaces to the MSTP topology.
3. Create VLANs mapped to MSTP instances tag interfaces to the VLANs.

(Step 1)

```
protocol spanning-tree mstp
  no disable
  name Tahiti
  revision 123
  MSTI 1 VLAN 100
  MSTI 2 VLAN 200,300
```

!

(Step 2)

```
interface GigabitEthernet 1/21
  no ip address
  switchport
  no shutdown
```

!

```
interface GigabitEthernet 1/31
  no ip address
  switchport
  no shutdown
```

!

(Step 3)

```
interface Vlan 100
  no ip address
  tagged GigabitEthernet 1/21,31
  no shutdown
```

!

```
interface Vlan 200
  no ip address
  tagged GigabitEthernet 1/21,31
```

```

    no shutdown
!
interface Vlan 300
    no ip address
    tagged GigabitEthernet 1/21,31
    no shutdown

```

Router 2 Running-Configuration

This example uses the following steps:

1. Enable MSTP globally and set the region name and revision map MSTP instances to the VLANs.
2. Assign Layer-2 interfaces to the MSTP topology.
3. Create VLANs mapped to MSTP instances tag interfaces to the VLANs.

(Step 1)

```

protocol spanning-tree mstp
    no disable
    name Tahiti
    revision 123
    MSTI 1 VLAN 100
    MSTI 2 VLAN 200,300
!

```

(Step 2)

```

interface GigabitEthernet 2/11
    no ip address
    switchport
    no shutdown
!
interface GigabitEthernet 2/31
    no ip address
    switchport
    no shutdown
!

```

(Step 3)

```

interface Vlan 100
    no ip address
    tagged GigabitEthernet 2/11,31
    no shutdown
!
interface Vlan 200
    no ip address
    tagged GigabitEthernet 2/11,31
    no shutdown
!
interface Vlan 300
    no ip address
    tagged GigabitEthernet 2/11,31
    no shutdown

```

Router 3 Running-Configuration

This example uses the following steps:

1. Enable MSTP globally and set the region name and revision map MSTP instances to the VLANs.
2. Assign Layer-2 interfaces to the MSTP topology.
3. Create VLANs mapped to MSTP instances tag interfaces to the VLANs.

(Step 1)

```

protocol spanning-tree mstp
    no disable

```

```

    name Tahiti
    revision 123
    MSTI 1 VLAN 100
    MSTI 2 VLAN 200,300
!
(Step 2)
interface GigabitEthernet 3/11
    no ip address
    switchport
    no shutdown
!
interface GigabitEthernet 3/21
    no ip address
    switchport
    no shutdown
!
(Step 3)
interface Vlan 100
    no ip address
    tagged GigabitEthernet 3/11,21
    no shutdown
!
interface Vlan 200
    no ip address
    tagged GigabitEthernet 3/11,21
    no shutdown
!
interface Vlan 300
    no ip address
    tagged GigabitEthernet 3/11,21
    no shutdown

```

SFTOS Example Running-Configuration

This example uses the following steps:

1. Enable MSTP globally and set the region name and revision map MSTP instances to the VLANs.
2. Assign Layer-2 interfaces to the MSTP topology.
3. Create VLANs mapped to MSTP instances tag interfaces to the VLANs.

```

(Step 1)
spanning-tree
spanning-tree configuration name Tahiti
spanning-tree configuration revision 123
spanning-tree MSTi instance 1
spanning-tree MSTi vlan 1 100
spanning-tree MSTi instance 2
spanning-tree MSTi vlan 2 200
spanning-tree MSTi vlan 2 300

```

```

(Step 2)
interface 1/0/31
    no shutdown
    spanning-tree port mode enable
    switchport protected 0
exit

interface 1/0/32
    no shutdown
    spanning-tree port mode enable
    switchport protected 0
exit

```

(Step 3)

```
interface vlan 100
  tagged 1/0/31
  tagged 1/0/32
exit
```

```
interface vlan 200
  tagged 1/0/31
  tagged 1/0/32
exit
```

```
interface vlan 300
  tagged 1/0/31
  tagged 1/0/32
exit
```

Debugging and Verifying MSTP Configurations

To debug and verify MSTP configuration, use the following commands.

- Display BPDUs.
EXEC Privilege mode

`debug spanning-tree mstp bpdu`
- Display MSTP-triggered topology change messages.

`debug spanning-tree mstp events`

Example of Viewing MSTP Configuration

Example of Viewing the Debug Log (Successful MSTP Configuration)

Example of Viewing the Debug Log (Unsuccessful MSTP Configuration)

To ensure all the necessary parameters match (region name, region version, and VLAN to instance mapping), examine your individual routers.

To show various portions of the MSTP configuration, use the `show spanning-tree mst` commands.

To view the overall MSTP configuration on the router, use the `show running-configuration spanning-tree mstp` in EXEC Privilege mode.

To monitor and verify that the MSTP configuration is connected and communicating as desired, use the `debug spanning-tree mstp bpdu` command.

Key items to look for in the debug report include:

- MSTP flags indicate communication received from the same region.
 - As shown in the following, the MSTP routers are located in the same region.
 - Does the debug log indicate that packets are coming from a “Different Region”? If so, one of the key parameters is not matching.
- MSTP Region Name and Revision.
 - The configured name and revisions must be identical among all the routers.

- Is the Region name blank? That may mean that a name was configured on one router and but was not configured or was configured differently on another router (spelling and capitalization counts).
- MSTP Instances.
 - To verify the VLAN to MSTP instance mapping, use the `show` commands.
 - Are there "extra" MSTP instances in the Sending or Received logs? This may mean that an additional MSTP instance was configured on one router but not the others.

```
Dell#show run spanning-tree mstp
!
protocol spanning-tree mstp
  name Tahiti
  revision 123
  MSTI 1 VLAN 100
  MSTI 2 VLAN 200,300
no disable
```

"Same Region," shown in bold in the following example shows that the MSTP routers are in a single region.

```
Dell#debug spanning-tree mstp bpdu
MSTP debug bpdu is ON
Dell#
4w0d4h : MSTP: Sending BPDU on Tengig 2/21 :
ProtId: 0, Ver: 3, Bpdu Type: MSTP, Flags 0x6e
CIST Root Bridge Id: 32768:0001.e806.953e, Ext Path Cost: 0
Regional Bridge Id: 32768:0001.e806.953e, CIST Port Id: 128:470
Msg Age: 0, Max Age: 20, Hello: 2, Fwd Delay: 15, Ver1 Len: 0, Ver3 Len: 96
Name: Tahiti, Rev: 123, Int Root Path Cost: 0
Rem Hops: 20, Bridge Id: 32768:0001.e806.953e
4w0d4h : INST 1: Flags: 0x6e, Reg Root: 32768:0001.e806.953e, Int Root Cost: 0
Brg/Port Prio: 32768/128, Rem Hops: 20
INST 2: Flags: 0x6e, Reg Root: 32768:0001.e806.953e, Int Root Cost: 0
Brg/Port Prio: 32768/128, Rem Hops: 20

4w0d4h : MSTP: Received BPDU on Tengig 2/21 :
ProtId: 0, Ver: 3, Bpdu Type: MSTP, Flags 0x78Same Region
CIST Root Bridge Id: 32768:0001.e806.953e, Ext Path Cost: 0
Regional Bridge Id: 32768:0001.e806.953e, CIST Port Id: 128:470
Msg Age: 0, Max Age: 20, Hello: 2, Fwd Delay: 15, Ver1 Len: 0, Ver3 Len: 96
Name: Tahiti, Rev: 123, Int Root Path Cost: 0
Rem Hops: 19, Bridge Id: 32768:0001.e8d5.cbbd
4w0d4h : INST 1: Flags: 0x78, Reg Root: 32768:0001.e806.953e, Int Root Cost: 0
      Brg/Port Prio: 32768/128, Rem Hops: 19
      INST 2: Flags: 0x78, Reg Root: 32768:0001.e806.953e, Int Root Cost: 0
      Brg/Port Prio: 32768/128, Rem Hops: 19
```

The bold line in the following example shows that the MSTP routers are in different regions and are not communicating with each other.

```
4w0d4h : MSTP: Received BPDU on TenGig 2/21 :
ProtId: 0, Ver: 3, Bpdu Type: MSTP, Flags 0x78Different Region
CIST Root Bridge Id: 32768:0001.e806.953e, Ext Path Cost: 0
Regional Bridge Id: 32768:0001.e806.953e, CIST Port Id: 128:470
Msg Age: 0, Max Age: 20, Hello: 2, Fwd Delay: 15, Ver1 Len: 0, Ver
Name: Tahiti, Rev: 123, Int Root Path Cost: 0
Rem Hops: 20, Bridge Id: 32768:0001.e8d5.cbbd
4w0d4h : INST 1: Flags: 0x70, Reg Root: 32768:0001.e8d5.cbbd, Int
      Brg/Port Prio: 32768/128, Rem Hops: 20
```

INST 2: Flags: 0x70, Reg Root: 32768:0001.e8d5.cbbd, Int Root Cost
Brg/Port Prio: 32768/128, Rem Hops: 20

Multicast Features

Multicast features are supported on the MXL switch platform.

The Dell Networking operating system (OS) supports the following multicast protocols:

- [PIM Sparse-Mode \(PIM-SM\)](#)
- [PIM Source-Specific Mode \(PIM-SSM\)](#)
- [Internet Group Management Protocol \(IGMP\)](#)
- [Multicast Source Discovery Protocol \(MSDP\)](#)

Enabling IP Multicast

Prior to enabling any multicast protocols, you must enable multicast routing.

- Enable multicast routing.
CONFIGURATION mode

```
ip multicast-routing
```

Implementation Information

Because protocol control traffic in the Networking OS is redirected using the MAC address, and multicast control traffic and multicast data traffic might map to the same MAC address, the Dell Networking OS might forward data traffic with certain MAC addresses to the CPU in addition to control traffic.

As the upper 5 bits of an IP Multicast address are dropped in the translation, 32 different multicast group IDs all map to the same Ethernet address. For example, 224.0.0.5 is a known IP address for open shortest path first (OSPF) that maps to the multicast MAC address 01:00:5e:00:00:05. However, 225.0.0.5, 226.0.0.5, and so on, map to the same multicast MAC address. The Layer 2 forwarding information base (FIB) alone cannot differentiate multicast control traffic multicast data traffic with the same address, so if you use IP address 225.0.0.5 for data traffic, both the multicast data and OSPF control traffic match the same entry and are forwarded to the CPU. Therefore, do not use well-known protocol multicast addresses for data transmission, such as the following.

Protocol	Ethernet Address
OSPF	01:00:5e:00:00:05
	01:00:5e:00:00:06

Protocol	Ethernet Address
RIP	01:00:5e:00:00:09
NTP	01:00:5e:00:01:01
VRRP	01:00:5e:00:00:12
PIM-SM	01:00:5e:00:00:0d


- The Dell Networking OS implementation of MTRACE is in accordance with IETF draft *draft-fenner-traceroute-ipm*.
- Multicast is not supported on secondary IP addresses.
- Egress L3 ACL is not applied to multicast data traffic if you enable multicast routing.

First Packet Forwarding for Lossless Multicast

Beginning with the Dell Networking OS version 8.3.1.0, all initial multicast packets are forwarded to receivers to achieve lossless multicast.

In previous versions, when the Dell Networking system is an RP, all initial packets are dropped until PIM creates an (S,G) entry. When the system is an RP and a Source DR, these initial packet drops represent a loss of native data, and when the system is an RP only, the initial packets drops represent a loss of register packets.

Both scenarios might be unacceptable depending on the multicast application. Beginning with the Dell Networking OS versions noted here, when the system is the RP, and has receivers for a group G, it forwards all initial multicast packets for the group based on the (*,G) entry rather than discarding them until the (S,G) entry is created, making Dell Networking systems suitable for applications sensitive to multicast packet loss.

 **NOTE:** When a source begins sending traffic, the Source DR forwards the initial packets to the RP as encapsulated registered packets. These packets are forwarded via the soft path at a maximum rate of 70 packets/second. Incoming packets beyond this rate are dropped.

Multicast Policies

The Dell Networking OS offers parallel multicast features for IPv4.

IPv4 Multicast Policies

The following sections describe IPv4 multicast policies.

- [Limiting the Number of Multicast Routes](#)
- [Preventing a Host from Joining a Group](#)
- [Rate Limiting IGMP Join Requests](#)
- [Preventing a PIM Router from Forming an Adjacency](#)
- [Preventing a Source from Registering with the RP](#)


- [Preventing a PIM Router from Processing a Join](#)

Limiting the Number of Multicast Routes

When the total number of multicast routes on a system limit is reached, the Dell Networking OS does not process any IGMP or multicast listener discovery protocol (MLD) joins to PIM — though it still processes leave messages — until the number of entries decreases below 95% of the limit.

When the limit falls below 95% after hitting the maximum, the system begins relearning route entries through IGMP, MLD, and MSDP.

- If the limit is increased after it is reached, join subsequent join requests are accepted. In this case, increase the limit by at least 10% for IGMP and MLD to resume.
- If the limit is decreased after it is reached, the system does not clear the existing sessions. Entries are cleared after a timeout (you may also clear entries using `clear ip mroute`).

 **NOTE:** The Dell Networking OS waits at least 30 seconds between stopping and starting IGMP join processing. You may experience this delay when manipulating the limit after it is reached.

When the multicast route limit is reached, the system displays the following:

```
3w1d13h: %RPM0-P:RP2 %PIM-3-PIM_TIB_LIMIT: PIM TIB limit reached. No new
routes will
be learnt until TIB level falls below low watermark.
3w1d13h: %RPM0-P:RP2 %PIM-3-PIM_TIB_LIMIT: PIM TIB below low watermark.
Route learning
will begin.
```

To limit the number of multicast routes, use the following command.


- Limit the total number of multicast routes on the system.

CONFIGURATION mode

```
ip multicast-limit
```


The range is from 1 to 50000.

The default is **15000**.

 **NOTE:** The IN-L3-McastFib CAM partition is used to store multicast routes and is a separate hardware limit that exists per port-pipe. Any software-configured limit may supersede by this hardware space limitation. The opposite is also true, the CAM partition might not be exhausted at the time the system-wide route limit the `ip multicast-limit` command sets is reached.

Preventing a Host from Joining a Group

You can prevent a host from joining a particular group by blocking specific IGMP reports. Create an extended access list containing the permissible source-group pairs.

 **NOTE:** For rules in IGMP access lists, *source* is the multicast source, not the source of the IGMP packet. For IGMPv2, use the keyword `any` for *source* (as shown in the following example), because IGMPv2 hosts do not know in advance who the source is for the group in which they are interested.

To apply the access list, use the following command.

- Apply the access list.

INTERFACE mode

```
ip igmp access-group access-list-name
```

Dell Networking OS Behavior: Do not enter the `ip igmp access-group` command before creating the access-list. If you do, after entering your first deny rule, the system clears multicast routing table and re-learns all groups, even those not covered by the rules in the access-list, because there is an implicit *deny all* rule at the end of all access-lists. Therefore, configuring an IGMP join request filter in this order might result in data loss. If you must enter the `ip igmp access-group` command before creating the access-list, prevent the system from clearing the routing table by entering a *permit any* rule with high sequence number before you enter any other rules.

In the following example, virtual local area network (VLAN) 400 is configured with an access list to permit only IGMP reports for group 239.0.0.1. Though Receiver 2 sends a membership report for groups 239.0.0.1 and 239.0.0.2, a multicast routing table entry is created only for group 239.0.0.1. VLAN 300 has no access list limiting Receiver 1, so both IGMP reports are accepted, and two corresponding entries are created in the routing table.

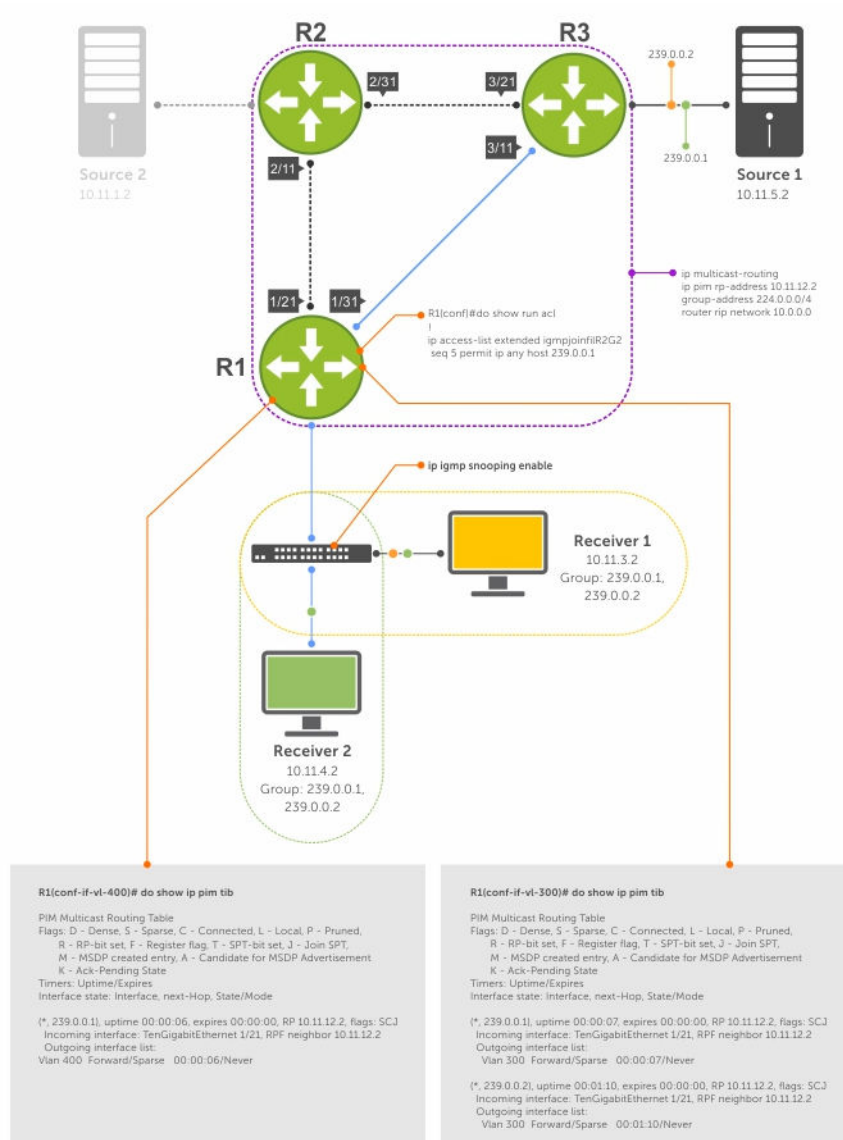


Figure 92. Preventing a Host from Joining a Group

Table 44. Preventing a Host from Joining a Group — Description

Location	Description
1/21	<ul style="list-style-type: none"> Interface GigabitEthernet 1/21 ip pim sparse-mode ip address 10.11.12.1/24 no shutdown
1/31	<ul style="list-style-type: none"> Interface GigabitEthernet 1/31 ip pim sparse-mode ip address 10.11.13.1/24

Location	Description
	<ul style="list-style-type: none"> no shutdown
2/1	<ul style="list-style-type: none"> Interface GigabitEthernet 2/1 ip pim sparse-mode ip address 10.11.1.1/24 no shutdown
2/11	<ul style="list-style-type: none"> Interface GigabitEthernet 2/11 ip pim sparse-mode ip address 10.11.12.2/24 no shutdown
2/31	<ul style="list-style-type: none"> Interface GigabitEthernet 2/31 ip pim sparse-mode ip address 10.11.23.1/24 no shutdown
3/1	<ul style="list-style-type: none"> Interface GigabitEthernet 3/1 ip pim sparse-mode ip address 10.11.5.1/24 no shutdown
3/11	<ul style="list-style-type: none"> Interface GigabitEthernet 3/11 ip pim sparse-mode ip address 10.11.13.2/24 no shutdown
3/21	<ul style="list-style-type: none"> Interface GigabitEthernet 3/21 ip pim sparse-mode ip address 10.11.23.2/24 no shutdown
Receiver 1	<ul style="list-style-type: none"> Interface VLAN 300 ip pim sparse-mode ip address 10.11.3.1/24 untagged GigabitEthernet 1/1 no shutdown
Receiver 2	<ul style="list-style-type: none"> Interface VLAN 400 ip pim sparse-mode ip address 10.11.4.1/24 untagged GigabitEthernet 1/2 ip igmp access-group igmpjoinfilR2G2 no shutdown

Rate Limiting IGMP Join Requests

If you expect a burst of IGMP Joins, protect the IGMP process from overload by limiting that rate at which new groups can be joined.

Hosts whose IGMP requests are denied will use the retry mechanism built-in to IGMP so that they're membership is delayed rather than permanently denied.

- Limit the rate at which new groups can be joined.

INTERFACE mode

```
ip igmp group-join-limit
```

To view the enable status of this feature, use the `show ip igmp interface` command from EXEC Privilege mode.

Preventing a PIM Router from Forming an Adjacency

To prevent a router from participating in PIM (for example, to configure stub multicast routing), use the following command.

- Prevent a router from participating in protocol independent multicast (PIM).

INTERFACE mode

```
ip pim neighbor-filter
```

Preventing a Source from Registering with the RP

To prevent the PIM source DR from sending register packets to RP for the specified multicast source and group, use the following command. If the source DR never sends register packets to the RP, no hosts can ever discover the source and create a shortest path tree (SPT) to it.

- Prevent a source from transmitting to a particular group.

CONFIGURATION mode

```
ip pim register-filter
```

In the following example, Source 1 and Source 2 are both transmitting packets for groups 239.0.0.1 and 239.0.0.2. R3 has a PIM register filter that only permits packets destined for group 239.0.0.2. An entry is created for group 239.0.0.1 in the routing table, but no outgoing interfaces are listed. R2 has no filter, so it is allowed to forward both groups. As a result, Receiver 1 receives only one transmission, while Receiver 2 receives duplicate transmissions.

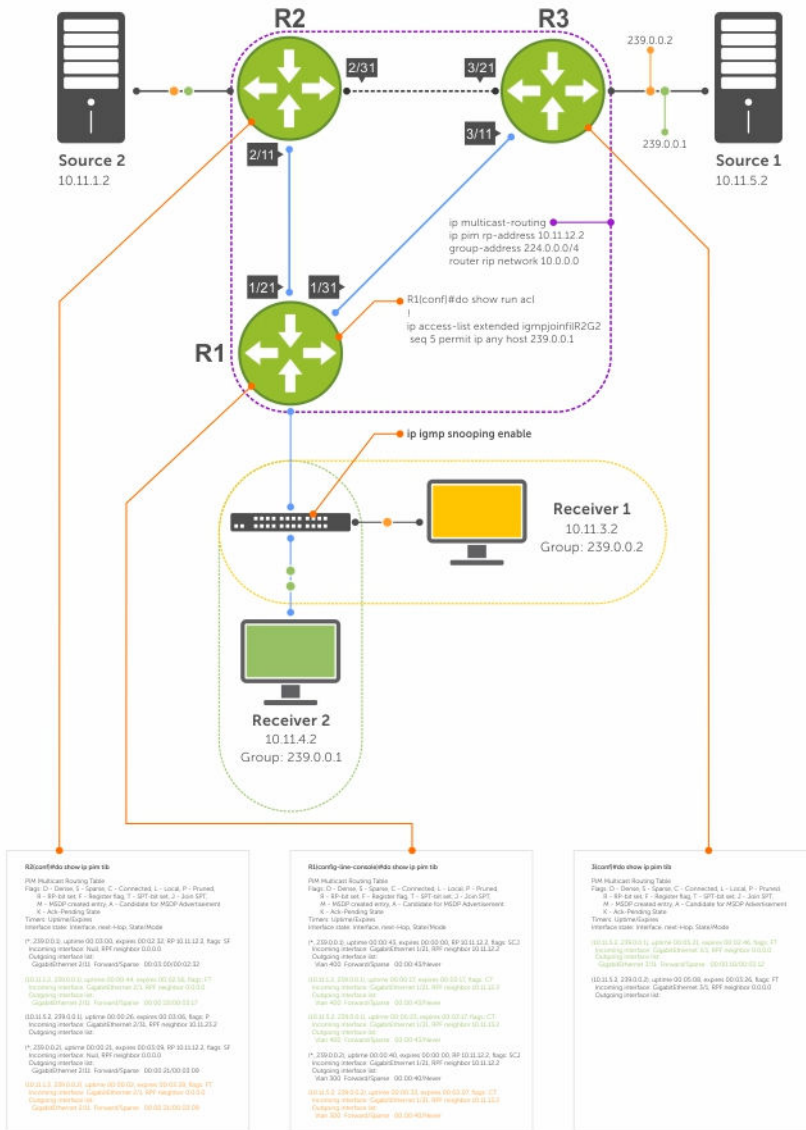


Figure 93. Preventing a Source from Transmitting to a Group


Table 45. Preventing a Source from Transmitting to a Group – Description

Location	Description
1/21	<ul style="list-style-type: none"> Interface GigabitEthernet 1/21 ip pim sparse-mode ip address 10.11.12.1/24 no shutdown
1/31	<ul style="list-style-type: none"> Interface GigabitEthernet 1/31 ip pim sparse-mode ip address 10.11.13.1/24

Location	Description
	<ul style="list-style-type: none"> no shutdown
2/1	<ul style="list-style-type: none"> Interface GigabitEthernet 2/1 ip pim sparse-mode ip address 10.11.1.1/24 no shutdown
2/11	<ul style="list-style-type: none"> Interface GigabitEthernet 2/11 ip pim sparse-mode ip address 10.11.12.2/24 no shutdown
2/31	<ul style="list-style-type: none"> Interface GigabitEthernet 2/31 ip pim sparse-mode ip address 10.11.23.1/24 no shutdown
3/1	<ul style="list-style-type: none"> Interface GigabitEthernet 3/1 ip pim sparse-mode ip address 10.11.5.1/24 no shutdown
3/11	<ul style="list-style-type: none"> Interface GigabitEthernet 3/11 ip pim sparse-mode ip address 10.11.13.2/24 no shutdown
3/21	<ul style="list-style-type: none"> Interface GigabitEthernet 3/21 ip pim sparse-mode ip address 10.11.23.2/24 no shutdown
Receiver 1	<ul style="list-style-type: none"> Interface VLAN 300 ip pim sparse-mode ip address 10.11.3.1/24 untagged GigabitEthernet 1/1 no shutdown
Receiver 2	<ul style="list-style-type: none"> Interface VLAN 400 ip pim sparse-mode ip address 10.11.4.1/24 untagged GigabitEthernet 1/2 no shutdown

Preventing a PIM Router from Processing a Join

To permit or deny PIM Join/Prune messages on an interface using an extended IP access list, use the following command.

 **NOTE:** Dell Networking recommends not using the `ip pim join-filter` command on an interface between a source and the RP router. Using this command in this scenario could cause problems with the PIM-SM source registration process resulting in excessive traffic being sent to the CPU of both the RP and PIM DR of the source.

Excessive traffic is generated when the join process from the RP back to the source is blocked due to a new source group being permitted in the join-filter. This results in the new source becoming stuck in registering on the DR and the continuous generation of UDP-encapsulated registration messages between the DR and RP routers which are being sent to the CPU.

- Prevent the PIM SM router from creating state based on multicast source and/ or group.


INTERFACE

```
ip pim join-filter
```


Open Shortest Path First (OSPFv2 and OSPFv3)

Open shortest path first (OSPFv2 for IPv4) and OSPF version 3 (OSPF for IPv6) are supported on the MXL switch platform.

This chapter provides a general description of OSPFv2 (OSPF for IPv4) and OSPFv3 (OSPF for IPv6) as supported in the Dell Networking operating system (OS).

 **NOTE:** The fundamental mechanisms of OSPF (flooding, DR election, area support, SPF calculations, and so on) are the same between OSPFv2 and OSPFv3. This chapter identifies and clarifies the differences between the two versions of OSPF. Except where identified, the information in this chapter applies to both protocol versions.

OSPF protocol standards are listed in the [Standards Compliance](#) chapter.

Protocol Overview

OSPF routing is a link-state routing protocol that calls for the sending of link-state advertisements (LSAs) to all other routers within the same autonomous system (AS) areas.

Information on attached interfaces, metrics used, and other variables is included in OSPF LSAs. As OSPF routers accumulate link-state information, they use the shortest path first (SPF) algorithm to calculate the shortest path to each node.

OSPF routers initially exchange HELLO messages to set up adjacencies with neighbor routers. The HELLO process is used to establish adjacencies between routers of the AS. It is not required that every router within the AS areas establish adjacencies. If two routers on the same subnet agree to become neighbors through the HELLO process, they begin to exchange network topology information in the form of LSAs.

In OSPFv2 neighbors on broadcast and NBMA links are identified by their interface addresses, while neighbors on other types of links are identified by RID.

Autonomous System (AS) Areas

OSPF operates in a type of hierarchy.

The largest entity within the hierarchy is the autonomous system (AS), which is a collection of networks under a common administration that share a common routing strategy. OSPF is an intra-AS (interior gateway) routing protocol, although it is capable of receiving routes from and sending routes to other ASs.

You can divide an AS into a number of areas, which are groups of contiguous networks and attached hosts. Routers with multiple interfaces can participate in multiple areas. These routers, called area border routers (ABRs), maintain separate databases for each area. Areas are a logical grouping of OSPF routers identified by an integer or dotted-decimal number.

Areas allow you to further organize your routers within in the AS. One or more areas are required within the AS. Areas are valuable in that they allow sub-networks to "hide" within the AS, thus minimizing the size of the routing tables on all routers. An area within the AS may not see the details of another area's topology. AS areas are known by their area number or the router's IP address.

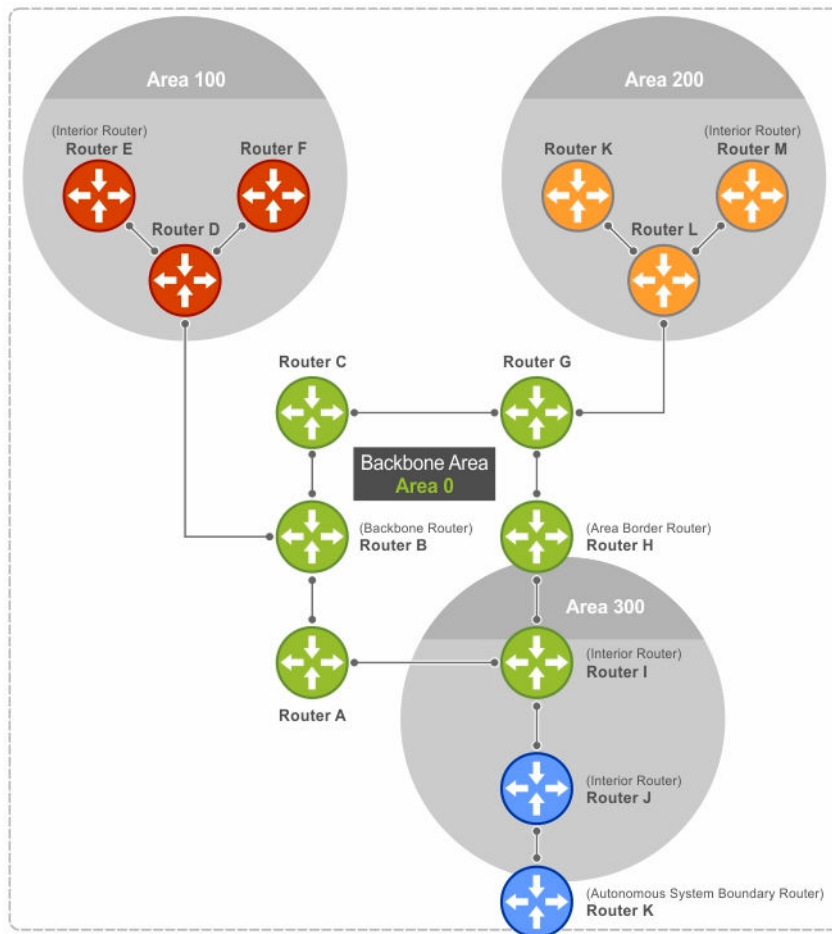


Figure 94. Autonomous System Areas

Area Types

The backbone of the network is Area 0. It is also called Area 0.0.0.0 and is the core of any AS.

All other areas must connect to Area 0. Areas can be defined in such a way that the backbone is not contiguous. In this case, backbone connectivity must be restored through virtual links. Virtual links are configured between any backbone routers that share a link to a non-backbone area and function as if they were direct links.

An OSPF backbone is responsible for distributing routing information between areas. It consists of all area border routers, networks not wholly contained in any area, and their attached routers.

The backbone is the only area with a default area number. All other areas can have their Area ID assigned in the configuration.

In the previous example, Routers A, B, C, G, H, and I are the Backbone.

- A stub area (SA) does not receive external route information, except for the default route. These areas do receive information from inter-area (IA) routes.
 - ✍ **NOTE:** Configure all routers within an assigned stub area as stubby, and not generate LSAs that do not apply. For example, a Type 5 LSA is intended for external areas and the Stubby area routers may not generate external LSAs. A virtual link cannot traverse stubby areas.
- A not-so-stubby area (NSSA) can import AS external route information and send it to the backbone. It cannot receive external AS information from the backbone or other areas. However, a virtual link can traverse it.
- Totally stubby areas are referred to as no summary areas in the Dell Networking OS.

Networks and Neighbors

As a link-state protocol, OSPF sends routing information to other OSPF routers concerning the state of the links between them. The state (up or down) of those links is important.

Routers that share a link become neighbors on that segment. OSPF uses the Hello protocol as a neighbor discovery and keep alive mechanism. After two routers are neighbors, they may proceed to exchange and synchronize their databases, which creates an adjacency.

Router Types

Router types are attributes of the OSPF process.

A given physical router may be a part of one or more OSPF processes. For example, a router connected to more than one area, receiving routing from a border gateway protocol (BGP) process connected to another AS acts as both an area border router and an autonomous system router.

Each router has a unique ID, written in decimal format (A.B.C.D). You do not have to associate the router ID with a valid IP address. However, to make troubleshooting easier, Dell Networking recommends that the router ID and the router's IP address reflect each other.

The following example shows different router designations.




Figure 95. OSPF Routing Examples

Backbone Router (BR)

A backbone router (BR) is part of the OSPF Backbone, Area 0.

This includes all ABRs. It can also include any routers that connect only to the backbone and another ABR, but are only part of Area 0, such as Router I in the previous example.

Area Border Router (ABR)

Within an AS, an area border router (ABR) connects one or more areas to the backbone.

The ABR keeps a copy of the link-state database for every area it connects to, so it may keep multiple copies of the link state database. An ABR takes information it has learned on one of its attached areas and can summarize it before sending it out on other areas it is connected to.

An ABR can connect to many areas in an AS, and is considered a member of each area it connects to.

Internal Router (IR)

The internal router (IR) has adjacencies with ONLY routers in the same area, as Router E, M, and I shown in the previous example.

Designated and Backup Designated Routers

OSPF elects a designated router (DR) and a backup designated router (BDR). Among other things, the DR is responsible for generating LSAs for the entire multiaccess network.

Designated routers allow a reduction in network traffic and in the size of the topological database.

- The DR maintains a complete topology table of the network and sends the updates to the other routers via multicast. All routers in an area form a slave/master relationship with the DR. Every time a router sends an update, the router sends it to the DR and BDR. The DR sends the update out to all other routers in the area.
- The BDR is the router that takes over if the DR fails.

Each router exchanges information with the DR and BDR. The DR and BDR relay the information to the other routers. On broadcast network segments, the number of OSPF packets is further reduced by the DR and BDR sending such OSPF updates to a multicast IP address that all OSPF routers on the network segment are listening on.

These router designations are not the same as the router IDs described earlier. The DRs and BDRs are configurable in the Dell Networking OS. If you do not define DR or BDR in the Dell Networking OS, the system assigns them. OSPF looks at the priority of the routers on the segment to determine which routers are the DR and BDR. The router with the highest priority is elected the DR. If there is a tie, the router with the higher router ID takes precedence. After the DR is elected, the BDR is elected the same way. A router with a router priority set to zero cannot become the DR or BDR.

Link-State Advertisements (LSAs)

A link-state advertisement (LSA) communicates the router's local routing topology to all other local routers in the same area.

The LSA types supported by Dell Networking are defined as follows:

- **Type 1: Router LSA** — The router lists links to other routers or networks in the same area. Type 1 LSAs are flooded across their own area only. The link-state ID of the Type 1 LSA is the originating router ID.
- **Type 2: Network LSA** — The DR in an area lists which routers are joined within the area. Type 2 LSAs are flooded across their own area only. The link-state ID of the Type 2 LSA is the IP interface address of the DR.
- **Type 3: Summary LSA (OSPFv2), Inter-Area-Prefix LSA (OSPFv3)** — An ABR takes information it has learned on one of its attached areas and can summarize it before sending it out on other areas it is connected to. The link-state ID of the Type 3 LSA is the destination network number.

- **Type 4: AS Border Router Summary LSA (OSPFv2), Inter-Area-Router LSA (OSPFv3)** — In some cases, Type 5 External LSAs are flooded to areas where the detailed next-hop information may not be available. An ABR floods the information for the router (for example, the ASBR where the Type 5 advertisement originated. The link-state ID for Type 4 LSAs is the router ID of the described ASBR).
- **Type 5: External LSA** — These LSAs contain information imported into OSPF from other routing processes. They are flooded to all areas, except stub areas. The link-state ID of the Type 5 LSA is the external network number.
- **Type 7: LSA** — Routers in an NSSA do not receive external LSAs from ABRs, but are allowed to send external routing information for redistribution. They use Type 7 LSAs to tell the ABRs about these external routes, which the ABR then translates to Type 5 external LSAs and floods as normal to the rest of the OSPF network.
- **Type 8: Link LSA (OSPFv3)** — This LSA carries the IPv6 address information of the local links.
- **Type 9: Link Local LSA (OSPFv2), Intra-Area-Prefix LSA (OSPFv3)** — For OSPFv2, this is a link-local "opaque" LSA as defined by RFC2370. For OSPFv3, this LSA carries the IPv6 prefixes of the router and network links.
- **Type 11 - Grace LSA (OSPFv3)** — For OSPFv3 only, this LSA is a link-local "opaque" LSA sent by a restarting OSPFv3 router during a graceful restart.

For all LSA types, there are 20-byte LSA headers. One of the fields of the LSA header is the link-state ID.

Each router link is defined as one of four types: type 1, 2, 3, or 4. The LSA includes a link ID field that identifies, by the network number and mask, the object this link connects to.

Depending on the type, the link ID has different meanings.

- 1: point-to-point connection to another router/neighbor router.
- 2: connection to a transit network IP address of the DR.
- 3: connection to a stub network IP network/subnet number.
- 4: virtual link neighboring router ID.

LSA Throttling

LSA throttling provides configurable interval timers to improve OSPF convergence times.

The default OSPF static timers (5 seconds for transmission, 1 second for acceptance) ensures sufficient time for sending and resending LSAs and for system acceptance of arriving LSAs. However, some networks may require reduced intervals for LSA transmission and acceptance. Throttling timers allow for this improved convergence times.

The LSA throttling timers are configured in milliseconds, with the interval time increasing exponentially until a maximum time has been reached. If the maximum time is reached, the system, the system continues to transmit at the max-interval until twice the max-interval time has passed. At that point, the system reverts to the start-interval timer and the cycle begins again.

When you configure the LSA throttle timers, syslog messages appear, indicating the interval times, as shown below for the transmit timer (45000ms) and arrival timer (1000ms).

```
Mar 15 09:46:00: %STKUNIT0-M:CP %OSPF-4-LSA_BACKOFF: OSPF Process 10,Router lsa
id
2.2.2.2 router-id 2.2.2.2 is backed off to transmit after 45000ms
Mar 15 09:46:06: %STKUNIT0-M:CP %OSPF-4-LSA_BACKOFF: OSPF Process 10,Router lsa
id
3.3.3.3 rtrid 3.3.3.3 received before 1000ms time
```

Router Priority and Cost

Router priority and cost is the method the system uses to "rate" the routers.

For example, if not assigned, the system selects the router with the highest priority as the DR. The second highest priority is the BDR.

- Priority is a numbered rating 0 to 255. The higher the number, the higher the priority.
- Cost is a numbered rating 1 to 65535. The higher the number, the greater the cost. The cost assigned reflects the cost should the router fail. When a router fails and the cost is assessed, a new priority number results.

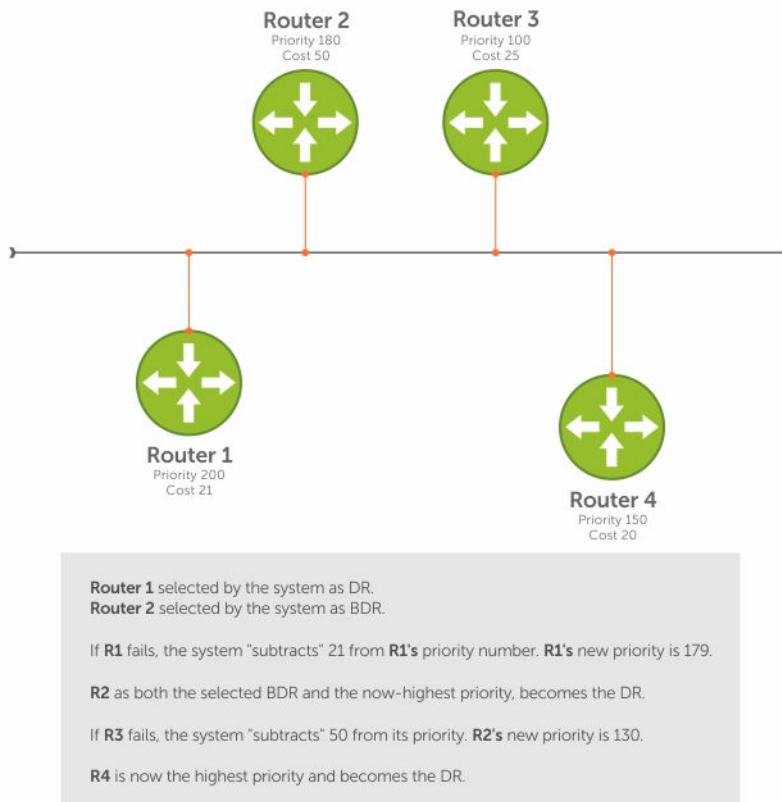


Figure 96. Priority and Cost Examples

OSPF with the Dell Networking OS

The Dell Networking OS supports up to 10,000 OSPF routes for OSPFv2. Within that 10,000 routes, you can designate up to 8,000 routes as external and up to 2,000 as inter/intra area routes.

The Dell Networking OS version 7.8.1.0 and later supports multiple OSPF processes (OSPF MP). The MXL switch supports up to 16 processes simultaneously.

On OSPFv3, the system supports only one process at a time for all platforms.

Prior to the Dell Networking OS version 7.8.1.0, the system supported one OSPFv2 and one OSPFv3 process ID per system. OSPFv2 and OSPFv3 can coexist but you must configure them individually.

The Dell Networking OS supports stub areas, totally stub (no summary) and not so stubby areas (NSSAs) and supports the following LSAs, as described earlier.

- Router (type 1)
- Network (type 2)
- Network Summary (type 3)
- AS Boundary (type 4)
- LSA(type 5)
- External LSA (type 7)
- Link LSA, OSPFv3 only (type 8)
- Opaque Link-Local (type 9)
- Grace LSA, OSPFv3 only (type 11)

Graceful Restart

Graceful restart for OSPFv2 and OSPFv3 are supported in Helper and Restart modes.

When a router goes down without a graceful restart, there is a possibility for loss of access to parts of the network due to ongoing network topology changes. Additionally, LSA flooding and reconvergence can cause substantial delays. It is, therefore, desirable that the network maintains a stable topology if it is possible for data flow to continue uninterrupted.

OSPF graceful restart understands that in a modern router, the control plane and data plane functionality are separate, restarting the control plane functionality (such as the failover of the active RPM to the backup in a redundant configuration), does not necessarily have to interrupt the forwarding of data packets. This behavior is supported because the forwarding tables previously computed by an active RPM have been downloaded into the forwarding information base (FIB) on the line cards (the data plane) and are still resident. For packets that have existing FIB/CAM entries, forwarding between ingress and egress ports/VLANs, and so on, can continue uninterrupted while the control plane OSPF process comes back to full functionality and rebuilds its routing tables.

To notify its helper neighbors that the restart process is beginning, when a router is attempting to restart gracefully, it originates the following link-local Grace LSAs:

- An OSPFv2 router sends Type 9 LSAs.
- An OSPFv3 router sends Type 11 LSAs.

Type 9 and 11 LSAs include a grace period, which is the time period an OSPF router advertises to adjacent neighbor routers as the time to wait for it to return to full control plane functionality. During the grace period, neighbor OSPFv2 /v3 interfaces save the LSAs from the restarting OSPF interface. Helper neighbor routers continue to announce the restarting router as fully adjacent, as long as the network topology remains unchanged. When the restarting router completes its restart, it flushes the Type 9 and 11 LSAs, notifying its neighbors that the restart is complete. This notification happens before the grace period expires.

Dell Networking routers support the following OSPF graceful restart functionality:

- Restarting role in which an enabled router performs its own graceful restart.
- Helper role in which the router's graceful restart function is to help a restarting neighbor router in its graceful restarts.
- Helper-reject role in which OSPF does not participate in the graceful restart of a neighbor.
- OSPFv2 supports *helper-only* and *restarting-only* roles. By default, both helper and restarting roles are enabled. OSPFv2 supports the helper-reject role globally on a router.
- OSPFv3 supports *helper-only* and *restarting-only* roles. The helper-only role is enabled by default. To enable the restarting role in addition to the helper-only role, configure a grace period. Reconfigure OSPFv3 graceful restart to a restarting-only role when you enable the helper-reject role on an interface. OSPFv3 supports the helper-reject role on a per-interface basis.
- Configuring helper-reject role on an OSPFv2 router or OSPFv3 interface enables the restarting-only role globally on the router or locally on the interface. In a helper-reject role, OSPF does not participate in the graceful restart of an adjacent OSPFv2/v3 router.
- If multiple OSPF interfaces provide communication between two routers, after you configure helper-reject on one interface, all other interfaces between the two routers behave as if they are in the helper-reject role.
- OSPFv2 and OSPFv3 support planned-only and/or unplanned-only restarts. The default is support for both planned and unplanned restarts.

A planned restart occurs when you enter the `redundancy force-failover rpm` command to force the primary RPM to switch to the backup RPM. During a planned restart, OSPF sends out a Grace LSA before the system switches over to the backup RPM.


An unplanned restart occurs when an unplanned event causes the active RPM to switch to the backup RPM, such as when an active process crashes, the active RPM is removed, or a power failure happens. During an unplanned restart, OSPF sends out a Grace LSA when the backup RPM comes online.

To display the configuration values for OSPF graceful restart, enter the `show run ospf` command for OSPFv2 and the `show run ospf` and `show ipv6 ospf database database-summary` commands for OSPFv3.

Fast Convergence (OSPFv2, IPv4 Only)

Fast convergence allows you to define the speeds at which LSAs are originated and accepted, and reduce OSPFv2 end-to-end convergence time.

The Dell Networking OS allows you to accept and originate LSAs as soon as they are available to speed up route information propagation.

 **NOTE:** The faster the convergence, the more frequent the route calculations and updates. This impacts CPU utilization and may impact adjacency stability in larger topologies.

Multi-Process OSPFv2 (IPv4 only)

Multi-process OSPF is supported on OSPFv2 with IPv4 only.

Multi-process OSPF allows multiple OSPFv2 processes on a single router. Multiple OSPFv2 processes allow for isolating routing domains, supporting multiple route policies and priorities in different domains, and creating smaller domains for easier management.

The MXL Switch supports up to 16 OSPFv2 processes.

Each OSPFv2 process has a unique process ID and must have an associated router ID. There must be an equal number of interfaces and must be in Layer-3 mode for the number of processes created. For

example, if you create five OSPFv2 processes on a system, there must be at least five interfaces assigned in Layer 3 mode.

Each OSPFv2 process is independent. If one process loses adjacency, the other processes continue to function.

Processing SNMP and Sending SNMP Traps

Though there are may be several OSPFv2 processes, only one process can process simple network management protocol (SNMP) requests and send SNMP traps.

The `mib-binding` command identifies one of the OSPFv2 processes as the process responsible for SNMP management. If you do not specify the `mib-binding` command, the first OSPFv2 process created manages the SNMP processes and traps.

RFC-2328 Compliant OSPF Flooding

In OSPF, flooding is the most resource-consuming task. The flooding algorithm described in RFC 2328 requires that OSPF flood LSAs on all interfaces, as governed by LSA's flooding scope (refer to Section 13 of the RFC.)

When multiple direct links connect two routers, the RFC 2328 flooding algorithm generates significant redundant information across all links.

By default, the Dell Networking OS implements an enhanced flooding procedure which dynamically and intelligently detects when to optimize flooding. Wherever possible, the OSPF task attempts to reduce flooding overhead by selectively flooding on a subset of the interfaces between two routers.

Enabling RFC-2328 Compliant OSPF Flooding

To enable OSPF flooding, use the following command.

When you enable this command, it configures the system to flood LSAs on all interfaces.

- Enable RFC 2328 flooding.

```
ROUTER OSPF mode
```

```
flood-2328
```

Example of Viewing the Debug Log for Flooding Behavior

Example of Verifying that RFC-2328–Compliant OSPF Flooding is Enabled

To confirm RFC 2328 flooding behavior, use the `debug ip ospf packet` command. The following example shows no change in the updated packets (shown in bold).

ACKs 2 (shown in bold) is printed only for ACK packets.

```
00:10:41 : OSPF(1000:00):
Rcv. v:2 t:5(LSAck) l:64 Acks 2 rid:2.2.2.2
    aid:1500 chk:0xdbee aut:0 auk: keyid:0 from:Vl 1000
        LSType:Type-5 AS External id:160.1.1.0 adv:6.1.0.0 seq:0x8000000c
        LSType:Type-5 AS External id:160.1.2.0 adv:6.1.0.0 seq:0x8000000c
00:10:41 : OSPF(1000:00):
Rcv. v:2 t:5(LSAck) l:64 Acks 2 rid:2.2.2.2
    aid:1500 chk:0xdbee aut:0 auk: keyid:0 from:Vl 100
        LSType:Type-5 AS External id:160.1.1.0 adv:6.1.0.0 seq:0x8000000c
        LSType:Type-5 AS External id:160.1.2.0 adv:6.1.0.0 seq:0x8000000c
```

```
00:10:41 : OSPF(1000:00):
Rcv. v:2 t:4 (LSUpd) l:100 rid:6.1.0.0
aid:0 chk:0xccbd aut:0 auk: keyid:0 from:Gi 10/21
Number of LSA:2
LSType:Type-5 AS External(5) Age:1 Seq:0x8000000c id:170.1.1.0 Adv:6.1.0.0
Netmask:255.255.255.0 fwd:0.0.0.0 E2, tos:0 metric:0
LSType:Type-5 AS External(5) Age:1 Seq:0x8000000c id:170.1.2.0 Adv:6.1.0.0
Netmask:255.255.255.0 fwd:0.0.0.0 E2, tos:0 metric:0
```

To confirm that you enabled RFC-2328-compliant OSPF flooding, use the `show ip ospf` command.

```
Dell#show ip ospf
Routing Process ospf 1 with ID 2.2.2.2
Supports only single TOS (TOS0) routes
It is an Autonomous System Boundary Router
It is Flooding according to RFC 2328
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Number of area in this router is 1, normal 0 stub 0 nssa 1
--More--
```

OSPF ACK Packing

The OSPF ACK packing feature bundles multiple LS acknowledgements in a single packet, significantly reducing the number of ACK packets transmitted when the number of LSAs increases.

This feature also enhances network utilization and reduces the number of small ACK packets sent to a neighboring router. OSPF ACK packing is enabled by default and non-configurable.

Setting OSPF Adjacency with Cisco Routers

To establish an OSPF adjacency between Dell Networking and Cisco routers, the hello interval and dead interval must be the same on both routers.

In Dell, the OSPF dead interval value is, by default, set to **40 seconds**, and is independent of the OSPF hello interval. Configuring a hello interval does not change the dead interval in Dell. In contrast, the OSPF dead interval on a Cisco router is, by default, four times as long as the hello interval. Changing the hello interval on the Cisco router automatically changes the dead interval.

For more information regarding this functionality or for assistance, go to <http://www.dell.com/support/my-support/>.

To ensure equal intervals between the routers, use the following command.

- Manually set the dead interval of the Dell Networking router to match the Cisco configuration.

```
INTERFACE mode

ip ospf dead-interval <x>
```

Example of Setting a Dead Interval

Example of Viewing Dead Interval Setting

In the following example, the dead interval is set at 4x the hello interval (shown in bold).

```
Dell(conf)#int tengig 2/2
Dell(conf-if-te-2/2)#ip ospf hello-interval 20
Dell(conf-if-te-2/2)#ip ospf dead-interval 80
```

```
Dell(conf-if-te-2/2)#
```

In the following example, the dead interval is set at 4x the hello interval (shown in bold).

```
Dell (conf-if-te-2/2)#ip ospf dead-interval 20
Dell (conf-if-te-2/2)#do show ip ospf int tengig 1/3
TenGigabitEthernet 2/2 is up, line protocol is up
  Internet Address 20.0.0.1/24, Area 0
  Process ID 10, Router ID 1.1.1.2, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 1.1.1.2, Interface address 30.0.0.1
  Backup Designated Router (ID) 1.1.1.1, Interface address 30.0.0.2
  Timer intervals configured, Hello 20, Dead 80, Wait 20, Retransmit 5
  Hello due in 00:00:04
  Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 1.1.1.1 (Backup Designated Router)
Dell (conf-if-gi-2/2)#
```


Configuration Information

The interfaces must be in Layer-3 mode (assigned an IP address) and enabled so that they can send and receive traffic. The OSPF process must know about these interfaces.

To make the OSPF process aware of these interfaces, they must be assigned to OSPF areas.

You must configure OSPF GLOBALLY on the system in CONFIGURATION mode.

OSPF features and functions are assigned to each router using the CONFIG-INTERFACE commands for each interface.

 **NOTE:** By default, OSPF is disabled.

Configuration Task List for OSPFv2 (OSPF for IPv4)

Open shortest path first version 2 (OSPF for IPv4) is supported on the MXL switch.

The following configuration tasks include two mandatory tasks and several optional tasks:

- [Enabling OSPFv2](#) (mandatory)
- [Enabling Multi-Process OSPF \(OSPFv2 IPv4 Only\)](#)
- [Assigning an OSPFv2 Area](#) (mandatory)
- [Enable OSPFv2 on Interfaces](#)
- [Configuring Stub Areas](#)
- [Configuring LSA Throttling Timers](#)
- [Enabling Passive Interfaces](#)
- [Enabling Fast-Convergence](#)
- [Changing OSPFv2 Parameters on Interfaces](#)
- [Enabling OSPFv2 Authentication](#)
- [Redistributing Routes](#)

- [Troubleshooting OSPFv2](#)

1. Configure a physical interface. Assign an IP address, physical or Loopback, to the interface to enable Layer 3 routing.
2. Enable OSPF globally. Assign network area and neighbors.
3. Add interfaces or configure other attributes.
4. Set the time interval between when the switch receives a topology change and starts a shortest path first (SPF) calculation.

Use `timers spf delay holdtime`

Example

```
Dell#
Dell#conf
Dell(conf)#router ospf 1
Dell(conf-router_ospf-1)#timer spf 2 5
Dell(conf-router_ospf-1)#
Dell(conf-router_ospf-1)#show config
!
router ospf 1
timers spf 2 5
Dell(conf-router_ospf-1)#
Dell(conf-router_ospf-1)#end
Dell#
```

For a complete list of the OSPF commands, refer to the *OSPF* section in the *Dell Networking OS Command Line Reference Guide* document.

Enabling OSPFv2

To enable Layer 3 routing, assign an IP address to an interface (physical or Loopback). By default, OSPF, similar to all routing protocols, is disabled.

You *must* configure at least one interface for Layer 3 before enabling OSPFv2 globally.

If implementing multi-process OSPF, create an equal number of Layer 3 enabled interfaces and OSPF process IDs. For example, if you create four OSPFv2 process IDs, you must have four interfaces with Layer 3 enabled.

1. Assign an IP address to an interface.

CONFIG-INTERFACE mode

```
ip address ip-address mask
```

The format is A.B.C.D/M.

If you are using a Loopback interface, refer to [Loopback Interfaces](#).

2. Enable the interface.

CONFIG-INTERFACE mode

```
no shutdown
```

3. Return to CONFIGURATION mode to enable the OSPFv2 process globally.

CONFIGURATION mode

```
router ospf process-id [vrf {vrf name}]
```

- *vrf name*: enter the keyword `VRF` and the instance name to tie the OSPF instance to the VRF. All network commands under this OSPF instance are later tied to the VRF instance.

The range is from 0 to 65535.

The OSPF process ID is the identifying number assigned to the OSPF process. The router ID is the IP address associated with the OSPF process.

After the OSPF process and the VRF are tied together, the OSPF process ID cannot be used again in the system.

If you try to enter an OSPF process ID, or if you try to enable more OSPF processes than available Layer 3 interfaces, prior to assigning an IP address to an interface and setting the `no shutdown` command, the following message displays:

```
Dell(conf)#router ospf 1
% Error: No router ID available..
```

Assigning a Router ID

In CONFIGURATION ROUTER OSPF mode, assign the router ID.

The router ID is not required to be the router's IP address. However, Dell Networking recommends using the IP address as the router ID for easier management and troubleshooting. Optional *process-id* commands are also described.

- Assign the router ID for the OSPFv2 process.

CONFIG-ROUTER-OSPF-id mode

```
router-id ip address
```

- Disable OSPF.

CONFIGURATION mode

```
no router ospf process-id
```

- Reset the OSPFv2 process.

EXEC Privilege mode

```
clear ip ospf process-id
```

- View the current OSPFv2 status.

EXEC mode

```
show ip ospf process-id
```

Example of Viewing the Current OSPFv2 Status

```
Dell#show ip ospf 55555
Routing Process ospf 55555 with ID 10.10.10.10
Supports only single TOS (TOS0) routes
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Number of area in this router is 0, normal 0 stub 0 nssa 0
Dell#
```

Enabling Multi-Process OSPF (OSPFv2, IPv4 Only)

Multi-process OSPF allows multiple OSPFv2 processes on a single router. The MXL switch supports up to 16 OSPFv2 processes.

When configuring a single OSPF process, follow the same steps previously described. Repeat them as often as necessary for the desired number of processes. After the process is created, all other configurations apply as usual.

1. Assign an IP address to an interface.

CONFIG-INTERFACE mode

```
ip address ip-address mask
```

Format: A.B.C.D/M.

If you are using a Loopback interface, refer to [Loopback Interfaces](#).

2. Enable the interface.

CONFIG-INTERFACE mode

```
no shutdown
```

3. Return to CONFIGURATION mode to enable the OSPFv2 process globally.

CONFIGURATION mode

```
router ospf process-id [vrf]
```

The range is from 0 to 65535.

After the OSPF process and the VRF are tied together, the OSPF process ID cannot be used again in the system.

If you try to enable more OSPF processes than available Layer 3 interfaces, the following message displays:

```
Dell(conf)#router ospf 1
% Error: No router ID available.
```

Assigning an OSPFv2 Area

After you enable OSPFv2, assign the interface to an OSPF area. Set up OSPF areas and enable OSPFv2 on an interface with the `network` command.

You must have at least one AS area: Area 0. This is the backbone area. If your OSPF network contains more than one area, configure a backbone area (Area ID 0.0.0.0). Any area besides Area 0 can have any number ID assigned to it.

The OSPFv2 process evaluates the `network` commands in the order they are configured. Assign the network address that is most explicit first to include all subnets of that address. For example, if you assign the network address 10.0.0.0 /8, you cannot assign the network address 10.1.0.0 /16 because it is already included in the first network address.

When configuring the `network` command, configure a network address and mask that is a superset of the IP subnet configured on the Layer-3 interface for OSPFv2 to use.

You can assign the area in the following step by a number or with an IP interface address.

- Enable OSPFv2 on an interface and assign a network address range to a specific OSPF area.
CONFIG-ROUTER-OSPF-id mode

```
network ip-address mask area area-id
```

The IP Address Format is A.B.C.D/M.


The area ID range is from 0 to 65535 or A.B.C.D/M.

Enable OSPFv2 on Interfaces

Enable and configure OSPFv2 on each interface (configure for Layer 3 protocol), and not shutdown.

You can also assign OSPFv2 to a Loopback interface as a virtual interface.

OSPF functions and features, such as MD5 Authentication, Grace Period, Authentication Wait Time, are assigned on a per interface basis.

 **NOTE:** If using features like MD5 Authentication, ensure all the neighboring routers are also configured for MD5.

In the following example, an IP address is assigned to an interface and an OSPFv2 area is defined that includes the IP address of a Layer 3 interface.

The first bold lines assign an IP address to a Layer 3 interface, and then `no shutdown` command ensures that the interface is UP.

The second bold line assigns the IP address of an interface to an area.

Example of Enabling OSPFv2 and Assigning an Area to an Interface

Example of Viewing Active Interfaces and Assigned Areas

Example of Viewing OSPF Status on a Loopback Interface

```
Dell#(conf)#int tengig 4/44
Dell(conf-if-te-4/44)#ip address 10.10.10.10/24
Dell(conf-if-te-4/44)#no shutdown
Dell(conf-if-te-4/44)#ex
Dell(conf)#router ospf 1
Dell(conf-router_ospf-1)#network 1.2.3.4/24 area 0
Dell(conf-router_ospf-1)#network 10.10.10.10/24 area 1
Dell(conf-router_ospf-1)#network 20.20.20.20/24 area 2
Dell(conf-router_ospf-1)#
Dell#
```

Dell Networking recommends using the interface IP addresses for the OSPFv2 router ID for easier management and troubleshooting.

To view the configuration, use the `show config` command in CONFIGURATION ROUTER OSPF mode.

OSPF, by default, sends hello packets out to all physical interfaces assigned an IP address that is a subset of a network on which OSPF is enabled.

To view currently active interfaces and the areas assigned to them, use the `show ip ospf interface` command.

```
Dell>show ip ospf 1 interface
```

```
TenGigabitEthernet 12/17 is up, line protocol is up
  Internet Address 10.2.2.1/24, Area 0.0.0.0
  Process ID 1, Router ID 11.1.2.1, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 11.1.2.1, Interface address 10.2.2.1
  Backup Designated Router (ID) 0.0.0.0, Interface address 0.0.0.0
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:04
  Neighbor Count is 0, Adjacent neighbor count is 0

TenGigabitEthernet 12/21 is up, line protocol is up
  Internet Address 10.2.3.1/24, Area 0.0.0.0
  Process ID 1, Router ID 11.1.2.1, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State BDR, Priority 1
  Designated Router (ID) 13.1.1.1, Interface address 10.2.3.2
  Backup Designated Router (ID) 11.1.2.1, Interface address 10.2.3.1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:05
  Neighbor Count is 1, Adjacent neighbor count is 1
  Adjacent with neighbor 13.1.1.1 (Designated Router)
Dell>
```

Loopback interfaces also help the OSPF process. OSPF picks the highest interface address as the router-id and a Loopback interface address has a higher precedence than other interface addresses.

```
Dell#show ip ospf 1 int
```

```
TenGigabitEthernet 13/23 is up, line protocol is up
  Internet Address 10.168.0.1/24, Area 0.0.0.1
  Process ID 1, Router ID 10.168.253.2, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DROTHER, Priority 1
  Designated Router (ID) 10.168.253.5, Interface address 10.168.0.4
  Backup Designated Router (ID) 192.168.253.3, Interface address 10.168.0.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:08
  Neighbor Count is 3, Adjacent neighbor count is 2
  Adjacent with neighbor 10.168.253.5 (Designated Router)
  Adjacent with neighbor 10.168.253.3 (Backup Designated Router)

Loopback 0 is up, line protocol is up
  Internet Address 10.168.253.2/32, Area 0.0.0.1
  Process ID 1, Router ID 10.168.253.2, Network Type LOOPBACK, Cost: 1
  Loopback interface is treated as a stub Host.
Dell#
```

Configuring Stub Areas

OSPF supports different types of LSAs to help reduce the amount of router processing within the areas. Type 5 LSAs are not flooded into stub areas; the ABR advertises a default route into the stub area to which it is attached. Stub area routers use the default route to reach external destinations.

To ensure connectivity in your OSPFv2 network, never configure the backbone area as a stub area.

To configure a stub area, use the following commands.

1. Review all areas after they were configured to determine which areas are NOT receiving type 5 LSAs.

EXEC Privilege mode

```
show ip ospf process-id database database-summary
```

2. Enter CONFIGURATION mode.

EXEC Privilege mode

```
configure
```

3. Enter ROUTER OSPF mode.

CONFIGURATION mode

```
router ospf process-id
```

Process ID is the ID assigned when configuring OSPFv2 globally.

4. Configure the area as a stub area.

CONFIG-ROUTER-OSPF-id mode

```
area area-id stub [no-summary]
```

Use the keywords `no-summary` to prevent transmission into the area of summary ASBR LSAs.

Area ID is the number or IP address assigned when creating the area.

Example of the `show ip ospf database database-summary` Command

To view which LSAs are transmitted, use the `show ip ospf database process-id database-summary` command in EXEC Privilege mode.

```
Dell#show ip ospf 34 database database-summary
```

```
OSPF Router with ID (10.1.2.100) (Process ID 34)

Area      ID Router Network S-Net S-ASBR Type-7 Subtotal
2.2.2.2   1      0      0      0      0      1
3.3.3.3   1      0      0      0      0      1
Dell#
```

To view information on areas, use the `show ip ospf process-id` command in EXEC Privilege mode.

Configuring LSA Throttling Timers

Configured LSA timers replace the standard transmit and acceptance times for LSAs.

The LSA throttling timers are configured in milliseconds, with the interval time increasing exponentially until a maximum time has been reached. If the maximum time is reached, the system continues to transmit at the max-interval. If the system is stable for twice the maximum interval time, the system reverts to the start-interval timer and the cycle begins again.

1. Specify the interval times for all LSA transmissions.

CONFIG-ROUTEROSPF- id mode

```
timers throttle lsa all {start-interval | hold-interval | max-interval}
```

- `start-interval`: set the minimum interval between the initial sending and resending the same LSA. The range is from 0 to 600,000 milliseconds.
 - `hold-interval`: set the next interval to send the same LSA. This interval is the time between sending the same LSA after the `start-interval` has been attempted. The range is from 1 to 600,000 milliseconds.
 - `max-interval`: set the maximum amount of time the system waits before sending the LSA. The range is from 1 to 600,000 milliseconds.
2. Specify the interval for LSA acceptance.
- CONFIG-ROUTEROSPF- id mode
- ```
timers throttle lsa arrival arrival-time
```
- `arrival-time`: set the interval between receiving the same LSA repeatedly, to allow sufficient time for the system to accept the LSA. The range is from 0 to 600,000 milliseconds.

## Enabling Passive Interfaces

A passive interface is one that does not send or receive routing information. Enabling passive interface suppresses routing updates on an interface. Although the passive interface does not send or receive routing updates, the network on that interface is still included in OSPF updates sent via other interfaces.

To suppress the interface's participation on an OSPF interface, use the following command. This command stops the router from sending updates on that interface.

- Specify whether all or some of the interfaces are passive.
- CONFIG-ROUTEROSPF- id mode

```
passive-interface {default | interface}
```

The default is enabled passive interfaces on ALL interfaces in the OSPF process.

Entering the physical interface type, slot, and number enables passive interface on only the identified interface.

- For a port channel, enter the keywords `port-channel` then a number from 1 to 128.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information (for example, `passive-interface ten 2/3`).
- For a VLAN, enter the keyword `vlan` then a number from 1 to 4094 (for example, `passive-interface vlan 2222`).
- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

The keyword `default` sets all interfaces on this OSPF process as passive.

To remove the passive interface from select interfaces, use the `no passive-interface interface` command while `passive interface default` is configured.

To enable both receiving and sending routing updates, use the `no passive-interface interface` command.

## Example of Viewing Passive Interfaces

When you configure a passive interface, the `show ip ospf process-id interface` command adds the words `passive interface` to indicate that the hello packets are not transmitted on that interface (shown in bold).

```
Dell#show ip ospf 34 int
```

```
TenGigabitEthernet 0/0 is up, line protocol is down
 Internet Address 10.1.2.100/24, Area 1.1.1.1
 Process ID 34, Router ID 10.1.2.100, Network Type BROADCAST, Cost: 10
 Transmit Delay is 1 sec, State DOWN, Priority 1
 Designated Router (ID) 10.1.2.100, Interface address 0.0.0.0
 Backup Designated Router (ID) 0.0.0.0, Interface address 0.0.0.0
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 Hello due in 13:39:46
 Neighbor Count is 0, Adjacent neighbor count is 0
```

```
TenGigabitEthernet 0/1 is up, line protocol is down
 Internet Address 10.1.3.100/24, Area 2.2.2.2
 Process ID 34, Router ID 10.1.2.100, Network Type BROADCAST, Cost: 10
 Transmit Delay is 1 sec, State DR, Priority 1
 Designated Router (ID) 10.1.2.100, Interface address 10.1.3.100
 Backup Designated Router (ID) 0.0.0.0, Interface address 0.0.0.0
 Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 No Hellos (Passive interface)
 Neighbor Count is 0, Adjacent neighbor count is 0
```

```
Loopback 45 is up, line protocol is up
 Internet Address 10.1.1.23/24, Area 2.2.2.2
 Process ID 34, Router ID 10.1.2.100, Network Type LOOPBACK, Cost: 1
```

## Enabling Fast-Convergence

The fast-convergence CLI sets the minimum origination and arrival LSA parameters to zero (0), allowing rapid route calculation.

When you disable fast-convergence, origination and arrival LSA parameters are set to 5 seconds and 1 second, respectively.

Setting the convergence parameter (from 1 to 4) indicates the actual convergence level. Each convergence setting adjusts the LSA parameters to zero, but the `fast-convergence` parameter setting allows for even finer tuning of the convergence speed. The higher the number, the faster the convergence.

To enable or disable fast-convergence, use the following command.


- Enable OSPF fast-convergence and specify the convergence level.  
CONFIG-ROUTEROSPF- id mode

```
fast-convergence {number}
```

The parameter range is from 1 to 4.

The higher the number, the faster the convergence.

When disabled, the parameter is set at 0.

 **NOTE:** A higher convergence level can result in occasional loss of OSPF adjacency. Generally, convergence level 1 meets most convergence requirements. Only select higher convergence levels following consultation with Dell Technical Support.

### Example of the `fast-converge` Command

### Example of Disabling Fast-Convergence

In the examples below, `Convergence Level` shows the fast-converge parameter setting and `Min LSA origination` shows the LSA parameters (shown in bold).

```
Dell(conf-router_ospf-1)#fast-converge 2
Dell(conf-router_ospf-1)#ex
Dell(conf)#ex
Dell#show ip ospf 1
Routing Process ospf 1 with ID 192.168.67.2
Supports only single TOS (TOS0) routes
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Convergence Level 2
Min LSA origination 0 secs, Min LSA arrival 0 secs
Number of area in this router is 0, normal 0 stub 0 nssa 0
Dell#
```

```
Dell#(conf-router_ospf-1)#no fast-converge
Dell#(conf-router_ospf-1)#ex
Dell#(conf)#ex
Dell##show ip ospf 1
Routing Process ospf 1 with ID 192.168.67.2
Supports only single TOS (TOS0) routes
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Convergence Level 0
Min LSA origination 0 msec, Min LSA arrival 1000 msec
Number of area in this router is 0, normal 0 stub 0 nssa 0
Dell#
```

## Changing OSPFv2 Parameters on Interfaces

In the Dell Networking OS, you can modify the OSPF settings on the interfaces. Some interface parameter values must be consistent across all interfaces to avoid routing errors. For example, set the same time interval for the hello packets on all routers in the OSPF network to prevent misconfiguration of OSPF neighbors.

To change OSPFv2 parameters on the interfaces, use any or all of the following commands.

- Change the cost associated with OSPF traffic on the interface.

CONFIG-INTERFACE mode

```
ip ospf cost
```

– *cost*: The range is from 1 to 65535 (the default depends on the interface speed).

- Change the time interval the router waits before declaring a neighbor dead.

CONFIG-INTERFACE mode

```
ip ospf dead-interval seconds
```

– *seconds*: the range is from 1 to 65535 (the default is **40 seconds**).

The dead interval must be four times the hello interval.

The dead interval must be the same on all routers in the OSPF network.

- Change the time interval between hello-packet transmission.

CONFIG-INTERFACE mode

```
ip ospf hello-interval seconds
```

- *seconds*: the range is from 1 to 65535 (the default is **10 seconds**).


The hello interval must be the same on all routers in the OSPF network.


- Use the MD5 algorithm to produce a message digest or key, which is sent instead of the key.

CONFIG-INTERFACE mode

```
ip ospf message-digest-key keyid md5 key
```

- *keyid*: the range is from 1 to 255.
- *Key*: a character string.

 **NOTE:** Be sure to write down or otherwise record the key. You cannot learn the key after it is configured. You must be careful when changing this key.

 **NOTE:** You can configure a maximum of six digest keys on an interface. Of the available six digest keys, the switches select the MD5 key that is common. The remaining MD5 keys are unused.

- Change the priority of the interface, which is used to determine the Designated Router for the OSPF broadcast network.

CONFIG-INTERFACE mode

```
ip ospf priority number
```

- *number*: the range is from 0 to 255 (the default is **1**).

- Change the retransmission interval between LSAs.

CONFIG-INTERFACE mode

```
ip ospf retransmit-interval seconds
```

- *seconds*: the range is from 1 to 65535 (the default is **5 seconds**).

The retransmit interval must be the same on all routers in the OSPF network.

- Change the wait period between link state update packets sent out the interface.

CONFIG-INTERFACE mode

```
ip ospf transmit-delay seconds
```

- *seconds*: the range is from 1 to 65535 (the default is **1 second**).

The transmit delay must be the same on all routers in the OSPF network.

### Example of Changing and Verifying the `cost` Parameter and Viewing Interface Status

To view interface configurations, use the `show config` command in CONFIGURATION INTERFACE mode.

To view interface status in the OSPF process, use the `show ip ospf interface` command in EXEC mode.

The bold lines in the example show the change on the interface. The change is reflected in the OSPF configuration.

```
Dell(conf-if)#ip ospf cost 45
Dell(conf-if)#show config
!
```

```
interface TenGigabitEthernet 0/0
 ip address 10.1.2.100 255.255.255.0
 no shutdown
 ip ospf cost 45
Dell(conf-if)#end
```

```
Dell#show ip ospf 34 interface
GigabitEthernet 0/0 is up, line protocol is up
Internet Address 10.1.2.100/24, Area 2.2.2.2
Process ID 34, Router ID 10.1.2.100, Network Type BROADCAST, Cost: 45
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router (ID) 10.1.2.100, Interface address 10.1.2.100
Backup Designated Router (ID) 10.1.2.100, Interface address 0.0.0.0
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
 Hello due in 00:00:06
Neighbor Count is 0, Adjacent neighbor count is 0
Dell#
```

## Enabling OSPFv2 Authentication

To enable or change various OSPF authentication parameters, use the following commands.

- Set a clear text authentication scheme on the interface.

CONFIG-INTERFACE mode

```
ip ospf authentication-key key
```

Configure a *key* that is a text string no longer than eight characters.

All neighboring routers must share password to exchange OSPF information.

- Set the authentication change wait time in seconds between 0 and 300 for the interface.

CONFIG-INTERFACE mode

```
ip ospf auth-change-wait-time seconds
```

This setting is the amount of time OSPF has available to change its interface authentication type.

During the *auth-change-wait-time*, OSPF sends out packets with both the new and old authentication schemes.

This transmission stops when the period ends.

The default is **0 seconds**.

## Enabling OSPFv2 Graceful Restart

Graceful restart is enabled for the global OSPF process.

For more information, refer to [Graceful Restart](#).

The Dell Networking implementation of OSPFv2 graceful restart enables you to specify:

- **grace period** — the length of time the graceful restart process can last before OSPF terminates it.
- **helper-reject neighbors** — the router ID of each restart router that does not receive assistance from the configured router.
- **mode** — the situation or situations that trigger a graceful restart.

- **role** — the role or roles the configured router can perform.

 **NOTE:** By default, OSPFv2 graceful restart is disabled.

To enable and configure OSPFv2 graceful restart, use the following commands.

1. Enable OSPFv2 graceful-restart globally and set the grace period.

```
CONFIG-ROUTEROSPF- id mode
```

```
graceful-restart grace-period seconds
```

The seconds range is from 40 and 3000.

This setting is the time that an OSPFv2 router's neighbors advertises it as fully adjacent, regardless of the synchronization state, during a graceful restart. OSPFv2 terminates this process when the grace period ends.

2. Enter the Router ID of the OSPFv2 helper router from which the router does not accept graceful restart assistance.

```
CONFIG-ROUTEROSPF- id mode
```

```
graceful-restart helper-reject router-id
```

- **Planned-only** — the OSPFv2 router supports graceful-restart for planned restarts only. A planned restart is when you manually enter a fail-over command to force the primary RPM over to the secondary RPM. During a planned restart, OSPF sends out a Grace LSA before the system switches over to the secondary RPM. OSPF also is notified that a planned restart is happening.
- **Unplanned-only** — the OSPFv2 router supports graceful-restart for only unplanned restarts. During an unplanned restart, OSPF sends out a Grace LSA after the secondary RPM comes online.

By default, OSPFv2 supports both planned and unplanned restarts. Selecting one or the other mode restricts OSPFv2 to the single selected mode.

3. Configure the graceful restart role or roles that this OSPFv2 router performs.

```
CONFIG-ROUTEROSPF- id mode
```

```
graceful-restart role [helper-only | restart-only]
```

The system supports the following options:

- **Helper-only:** the OSPFv2 router supports graceful-restart only as a helper router.
- **Restart-only:** the OSPFv2 router supports graceful-restart only during unplanned restarts.

By default, OSPFv2 supports both restarting and helper roles. Selecting one or the other role restricts OSPFv2 to the single selected role.

To disable OSPFv2 graceful-restart after you have enabled it, use the `no graceful-restart grace-period` command in CONFIG-ROUTEROSPF- id mode. The command returns OSPF graceful-restart to its default state.

For more information about OSPF graceful restart, refer to the *Dell Networking OS Command Line Reference Guide*.

#### Example of the `show run ospf` Command

When you configure a graceful restart on an OSPFv2 router, the `show run ospf` command displays information similar to the following.



```
Dell#show run ospf
!
router ospf 1
 graceful-restart grace-period 300
 graceful-restart role helper-only
 graceful-restart mode unplanned-only
 graceful-restart helper-reject 10.1.1.1
 graceful-restart helper-reject 20.1.1.1
 network 10.0.2.0/24 area 0
Dell#
```

## Creating Filter Routes

To filter routes, use prefix lists. OSPF applies prefix lists to incoming or outgoing routes. Incoming routes must meet the conditions of the prefix lists. If they do not, OSPF does not add the route to the routing table. Configure the prefix list in CONFIGURATION PREFIX LIST mode prior to assigning it to the OSPF process.

- Create a prefix list and assign it a unique name.

CONFIGURATION mode

```
ip prefix-list prefix-name
```

You are in PREFIX LIST mode.

- Create a prefix list with a sequence number and a deny or permit action.

CONFIG- PREFIX LIST mode

```
seq sequence-number {deny |permit} ip-prefix [ge min-prefix-length] [le max-prefix-length]
```

The optional parameters are:

- ge *min-prefix-length*: is the minimum prefix length to match (from 0 to 32).
- le *max-prefix-length*: is the maximum prefix length to match (from 0 to 32).

For configuration information about prefix lists, refer to [Access Control Lists \(ACLs\)](#).

## Applying Prefix Lists

To apply prefix lists to incoming or outgoing OSPF routes, use the following commands.

- Apply a configured prefix list to incoming OSPF routes.

CONFIG-ROUTEROSPF-id mode

```
distribute-list prefix-list-name in [interface]
```


- Assign a configured prefix list to outgoing OSPF routes.

CONFIG-ROUTEROSPF-id

```
distribute-list prefix-list-name out [connected | ospf <process if> | rip | static]
```

## Redistributing Routes

You can add routes from other routing instances or protocols to the OSPF process. With the `redistribute` command, you can include RIP, static, or directly connected routes in the OSPF process.

 **NOTE:** Do not route iBGP routes to OSPF unless there are route-maps associated with the OSPF redistribution.

To redistribute routes, use the following command.

- Specify which routes are redistributed into OSPF process.

CONFIG-ROUTEROSPF-id mode

```
redistribute {bgp | connected | rip | ospf <process id> | static} [metric
metric-value | metric-type type-value] [route-map map-name] [tag tag-value]
```

Configure the following required and optional parameters:

- `bgp`, `connected`, `ospf`, `rip`, `static`: enter one of the keywords to redistribute those routes.
- `metric metric-value`: the range is from 0 to 4294967295.
- `metric-type metric-type`: 1 for OSPF external route type 1. 2 for OSPF external route type 2.
- `route-map map-name`: enter a name of a configured route map.
- `tag tag-value`: the range is from 0 to 4294967295.

## Example of Viewing OSPF Configuration after Redistributing Routes


To view the current OSPF configuration, use the `show running-config ospf` command in EXEC mode or the `show config` command in ROUTER OSPF mode.

```
Dell(conf-router_ospf)#show config
!
router ospf 34
 network 10.1.2.32 0.0.0.255 area 2.2.2.2
 network 10.1.3.24 0.0.0.255 area 3.3.3.3
 distribute-list dilling in
Dell(conf-router_ospf)#
```

## Troubleshooting OSPFv2

The Dell Networking OS has several tools to make troubleshooting easier.

Be sure to check the following, as these questions represent typical issues that interrupt an OSPFv2 process.

 **NOTE:** The following is not a comprehensive list, just some examples of typical troubleshooting checks.

- Have you enabled OSPF globally?
- Is the OSPF process active on the interface?
- Are adjacencies established correctly?
- Are the interfaces configured for Layer 3 correctly?
- Is the router in the correct area type?
- Have the routes been included in the OSPF database?

- Have the OSPF routes been included in the routing table (not just the OSPF database)?

Some useful troubleshooting commands are:

- `show interfaces`
- `show protocols`
- `debug IP OSPF events and/or packets`
- `show neighbors`
- `show routes`

To help troubleshoot OSPFv2, use the following commands.

- View the summary of all OSPF process IDs enabled on the router.

EXEC Privilege mode

```
show running-config ospf
```

- View the summary information of the IP routes.

EXEC Privilege mode

```
show ip route summary
```

- View the summary information for the OSPF database.

EXEC Privilege mode

```
show ip ospf database
```

- View the configuration of OSPF neighbors connected to the local router.

EXEC Privilege mode

```
show ip ospf neighbor
```

- View the LSAs currently in the queue.

EXEC Privilege mode

```
show ip ospf timers rate-limit
```

- View debug messages.

EXEC Privilege mode

```
debug ip ospf process-id [event | packet | spf | database-timers rate-limit]
```

To view debug messages for a specific OSPF process ID, use the `debug ip ospf process-id` command.

If you do not enter a process ID, the command applies to the first OSPF process.

To view debug messages for a specific operation, enter one of the optional keywords:

- `event`: view OSPF event messages.
- `packet`: view OSPF packet information.
- `spf`: view SPF information.
- `database-timers rate-limit`: view the LSAs currently in the queue.

### Example of Viewing OSPF Configuration

```
Dell#show run ospf
!
```

```

router ospf 3
!
router ospf 4
 router-id 4.4.4.4
 network 4.4.4.0/28 area 1
!
router ospf 5
!
router ospf 6
!
router ospf 7
 mib-binding
!
router ospf 8
!
router ospf 90
 area 2 virtual-link 4.4.4.4
 area 2 virtual-link 90.90.90.90 retransmit-interval 300
!
ipv6 router ospf 999
 default-information originate always
 router-id 10.10.10.10
Dell#

```

## Sample Configurations for OSPFv2

The following configurations are examples for enabling OSPFv2.

These examples are not comprehensive directions. They are intended to give you some guidance with typical configurations.

You can copy and paste from these examples to your CLI. To support your own IP addresses, interfaces, names, and so on, be sure that you make the necessary changes.

### Basic OSPFv2 Router Topology

The following illustration is a sample basic OSPFv2 topology.

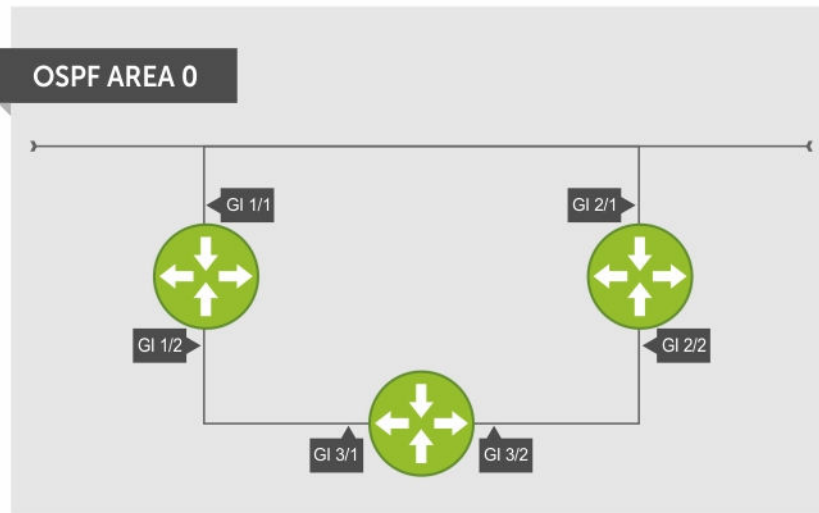


Figure 97. Basic Topology and CLI Commands for OSPFv2

### **OSPF Area 0 – G1/1 and 1/2**

```
router ospf 11111
 network 10.0.11.0/24 area 0
 network 10.0.12.0/24 area 0
 network 192.168.100.0/24 area 0
!
interface GigabitEthernet 1/1
 ip address 10.1.11.1/24
 no shutdown
!
interface GigabitEthernet 1/2
 ip address 10.2.12.2/24
 no shutdown
!
interface Loopback 10
 ip address 192.168.100.100/24
 no shutdown
```

### **OSPF Area 0 – G1/3/1 and 3/2**

```
router ospf 33333
 network 192.168.100.0/24 area 0
 network 10.0.13.0/24 area 0
 network 10.0.23.0/24 area 0
!
interface Loopback 30
 ip address 192.168.100.100/24
 no shutdown
!
interface GigabitEthernet 3/1
 ip address 10.1.13.3/24
 no shutdown
!
interface GigabitEthernet 3/2
 ip address 10.2.13.3/24
 no shutdown
```

### **OSPF Area 0 – G1/2/1 and 2/2**

```
router ospf 22222
 network 192.168.100.0/24 area 0
 network 10.2.21.0/24 area 0
 network 10.2.22.0/24 area 0
!
interface Loopback 20
 ip address 192.168.100.20/24
 no shutdown
!
interface GigabitEthernet 2/1
 ip address 10.2.21.2/24
 no shutdown
!
interface GigabitEthernet 2/2
 ip address 10.2.22.2/24
 no shutdown
```


## Configuration Task List for OSPFv3 (OSPF for IPv6)

The configuration options of OSPFv3 are the same as those options for OSPFv2, but you may configure OSPFv3 with differently labeled commands.

Specify process IDs and areas and include interfaces and addresses in the process. Define areas as stub or totally stubby.

The interfaces must be in IPv6 Layer-3 mode (assigned an IPv6 IP address) and enabled so that they can send and receive traffic. The OSPF process must know about these interfaces. To make the OSPF process aware of these interfaces, assign them to OSPF areas.

The OSPFv3 `ipv6 ospf area` command enables OSPFv3 on the interface and places the interface in an area. With OSPFv2, two commands are required to accomplish the same tasks – the `router ospf` command to create the OSPF process, then the `network area` command to enable OSPF on an interface.

 **NOTE:** The OSPFv2 `network area` command enables OSPF on multiple interfaces with the single command. Use the OSPFv3 `ipv6 ospf area` command on each interface that runs OSPFv3.

All IPv6 addresses on an interface are included in the OSPFv3 process that is created on the interface.

Enable OSPFv3 for IPv6 by specifying an OSPF process ID and an area in INTERFACE mode. If you have not created an OSPFv3 process, it is created automatically. All IPv6 addresses configured on the interface are included in the specified OSPF process.

Set the time interval between when the switch receives a topology change and starts a shortest path first (SPF) calculation

Use `timers spf delay holdtime`

### Example

```
Dell#
Dell#conf
Dell(conf)#ipv6 router ospf 1
Dell(conf-ipv6-router_ospf)#timer spf 2 5
Dell(conf-ipv6-router_ospf)#
Dell(conf-ipv6-router_ospf)#show config
!
ipv6 router ospf 1
timers spf 2 5
Dell(conf-ipv6-router_ospf)#
Dell(conf-ipv6-router_ospf)#end
Dell#
```

 **NOTE:** IPv6 and OSPFv3 do not support Multi-Process OSPF. You can only enable a single OSPFv3 process.

## Enabling IPv6 Unicast Routing

To enable IPv6 unicast routing, use the following command.

- Enable IPv6 unicast routing globally.

CONFIGURATION mode

```
ipv6 unicast routing
```

## Assigning IPv6 Addresses on an Interface

To assign IPv6 addresses to an interface, use the following commands.

1. Assign an IPv6 address to the interface.

CONF-INT-type slot/port mode

```
ipv6 address ipv6 address
```

IPv6 addresses are normally written as eight groups of four hexadecimal digits; separate each group by a colon (:).

The format is A:B:C::F/128.

2. Bring up the interface.


CONF-INT-type slot/port mode

```
no shutdown
```

## Assigning Area ID on an Interface

To assign the OSPFv3 process to an interface, use the following command.

The `ipv6 ospf area` command enables OSPFv3 on an interface and places the interface in the specified area. Additionally, the command creates the OSPFv3 process with ID on the router. OSPFv2 requires two commands to accomplish the same tasks – the `router ospf` command to create the OSPF process, then the `network area` command to enable OSPFv2 on an interface.

 **NOTE:** The OSPFv2 `network area` command enables OSPFv2 on multiple interfaces with the single command. Use the OSPFv3 `ipv6 ospf area` command on each interface that runs OSPFv3.

- Assign the OSPFv3 process and an OSPFv3 area to this interface.

CONF-INT-type slot/port mode

```
ipv6 ospf process-id area area-id
```

– *process-id*: the process ID number assigned.

– *area-id*: the area ID for this interface.

## Assigning OSPFv3 Process ID and Router ID Globally

To assign, disable, or reset OSPFv3 globally, use the following commands.

- Enable the OSPFv3 process globally and enter OSPFv3 mode.  
CONFIGURATION mode

```
ipv6 router ospf {process ID}
```


The range is from 0 to 65535.

- Assign the router ID for this OSPFv3 process.  
CONF-IPV6-ROUTER-OSPF mode

```
router-id {number}
```

– *number*: the IPv4 address.

The format is A.B.C.D.

 **NOTE:** Enter the router-id for an OSPFv3 router as an IPv4 IP address.

- Disable OSPF.  
CONFIGURATION mode

```
no ipv6 router ospf process-id
```

- Reset the OSPFv3 process.  
EXEC Privilege mode

```
clear ipv6 ospf process
```

## Configuring Stub Areas

To configure IPv6 stub areas, use the following command.

- Configure the area as a stub area.  
CONF-IPV6-ROUTER-OSPF mode

```
area area-id stub [no-summary]
```

- *no-summary*: use these keywords to prevent transmission in to the area of summary ASBR LSAs.
- *Area ID*: a number or IP address assigned when creating the area. You can represent the area ID as a number from 0 to 65536 if you assign a dotted decimal format rather than an IP address.

## Configuring Passive-Interface

To suppress the interface's participation on an OSPFv3 interface, use the following command. This command stops the router from sending updates on that interface.

- Specify whether some or all some of the interfaces are passive.  
CONF-IPV6-ROUTER-OSPF mode



```
passive-interface {type slot/port}
```

Interface: identifies the specific interface that is passive.

- For a Gigabit Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information (for example, `passive-interface gi 2/1`).
- For a port channel, enter the keywords `port-channel` then a number from 1 to 255 (for example, `passive-interface po 100`).
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information (for example, `passive-interface ten 2/3`).
- For a 40-Gigabit Ethernet interface, enter the keyword `FortyGigabitEthernet` then the slot/port information.
- For a VLAN, enter the keyword `vlan` then a number from 1 to 4094 (for example, `passive-interface vlan 2222`).

To enable both receiving and sending routing updates, use the `no passive-interface interface` command.

To indicate that hello packets are not transmitted on that interface, when you configure a passive interface, the `show ipv6 ospf interface` command adds the words `passive interface`.

## Redistributing Routes

You can add routes from other routing instances or protocols to the OSPFv3 process.

With the `redistribute` command, you can include RIP, static, or directly connected routes in the OSPF process. Route redistribution is also supported between OSPF Routing process IDs.

To add redistributing routes, use the following command.

- Specify which routes are redistributed into the OSPF process.

CONF-IPV6-ROUTER-OSPF mode

```
redistribute {bgp | connected | static} [metric metric-value | metric-type
type-value] [route-map map-name] [tag tag-value]
```

Configure the following required and optional parameters:

- `bgp | connected | static`: enter one of the keywords to redistribute those routes.
- `metric metric-value`: The range is from 0 to 4294967295.
- `metric-type metric-type`: enter 1 for OSPFv3 external route type 1 OR 2 for OSPFv3 external route type 2.
- `route-map map-name`: enter a name of a configured route map.
- `tag tag-value`: The range is from 0 to 4294967295.

## Configuring a Default Route

To generate a default external route into the OSPFv3 routing domain, configure the Dell Networking OS. To specify the information for the default route, use the following command.

- Specify the information for the default route.

CONF-IPV6-ROUTER-OSPF mode

```
default-information originate [always [metric metric-value] [metric-type
type-value]] [route-map map-name]
```

Configure the following required and optional parameters:

- `always`: indicate that default route information is always advertised.
- `metric metric-value`: The range is from 0 to 4294967295.
- `metric-type metric-type`: enter 1 for OSPFv3 external route type 1 OR 2 for OSPFv3 external route type 2.
- `route-map map-name`: enter a name of a configured route map.

## Enabling OSPFv3 Graceful Restart


Graceful restart for OSPFv3 is supported on the MXL switch.

For more information about graceful restart, refer to [Graceful Restart](#).

By default, OSPFv3 graceful restart is disabled and functions only in a helper role to help restarting neighbor routers in their graceful restarts when it receives a Grace LSA.

To enable OSPFv3 graceful restart, enter the `ipv6 router ospf process-id` command to enter OSPFv3 configuration mode. Then configure a grace period using the `graceful-restart grace-period` command. The grace period is the time that the OSPFv3 neighbors continue to advertise the restarting router as though it is fully adjacent. When you enable graceful restart (restarting role), an OSPFv3 restarting expects its OSPFv3 neighbors to help when it restarts by not advertising the broken link.

When you enable the helper-reject role on an interface using the `ipv6 ospf graceful-restart helper-reject` command, you reconfigure OSPFv3 graceful restart to function in a *restarting-only* role. OSPFv3 does not participate in the graceful restart of a neighbor.

 **NOTE:** Enter the `ipv6 ospf graceful-restart helper-reject` command in Interface configuration mode.

- Enable OSPFv3 graceful restart globally by setting the grace period (in seconds).  
CONF-IPV6-ROUTER-OSPF mode

```
graceful-restart grace-period seconds
```

The valid values are from 40 to 1800 seconds.

- Configure an OSPFv3 interface to not act on the Grace LSAs that it receives from a restarting OSPFv3 neighbor.

INTERFACE mode

```
ipv6 ospf graceful-restart helper-reject
```

- Specify the operating mode and type of events that trigger a graceful restart.  
CONF-IPV6-ROUTER-OSPF mode

```
graceful-restart mode [planned-only | unplanned-only]
```

- `Planned-only`: the OSPFv3 router supports graceful restart only for planned restarts. A planned restart is when you manually enter a `redundancy force-failover rpm` command to force the primary RPM over to the secondary RPM. During a planned restart, OSPFv3 sends out a Grace LSA

before the system switches over to the secondary RPM. OSPFv3 is notified that a planned restart is happening.

- **Unplanned-only**: the OSPFv3 router supports graceful-restart only for unplanned restarts. During an unplanned restart, OSPFv3 sends out a Grace LSA once the secondary RPM comes online.

The default is both planned and unplanned restarts trigger an OSPFv3 graceful restart. Selecting one or the other mode restricts OSPFv3 to the single selected mode.

- Disable OSPFv3 graceful-restart.  
CONF-IPV6-ROUTER-OSPF mode

```
no graceful-restart grace-period
```

## Displaying Graceful Restart

To display information on the use and configuration of OSPFv3 graceful restart, enter any of the following commands.

- Display the graceful-restart configuration for OSPFv2 and OSPFv3 (shown in the following example).  
EXEC Privilege mode

```
show run ospf
```

- Display the Type-11 Grace LSAs sent and received on an OSPFv3 router (shown in the following example).  
EXEC Privilege mode

```
show ipv6 ospf database grace-lsa
```

- Display the currently configured OSPFv3 parameters for graceful restart (shown in the following example).  
EXEC Privilege mode

```
show ipv6 ospf database database-summary
```

### Example of the `show run ospf` Command

### Example of the `show ipv6 ospf database database-summary` Command

### Example of the `show ipv6 ospf database grace-lsa` Command

```
Dell#show run ospf
!
router ospf 1
 router-id 200.1.1.1
 log-adjacency-changes
 graceful-restart grace-period 180
 network 20.1.1.0/24 area 0
 network 30.1.1.0/24 area 0
!
ipv6 router ospf 1
 log-adjacency-changes
 graceful-restart grace-period 180

Dell#show ipv6 ospf database database-summary
!
OSPFv3 Router with ID (200.1.1.1) (Process ID 1)
```

```

Process 1 database summary
Type Count/Status
Oper Status 1
Admin Status 1
Area Bdr Rtr Status 0
AS Bdr Rtr Status 1
AS Scope LSA Count 0
AS Scope LSA Cksum sum 0
Originate New LSAS 73
Rx New LSAS 114085
Ext LSA Count 0
Rte Max Eq Cost Paths 5
GR grace-period 180
GR mode planned and unplanned

```

```

Area 0 database summary
Type Count/Status
Brd Rtr Count 2
AS Bdr Rtr Count 2
LSA count 12010
Summary LSAs 1
Rtr LSA Count 4
Net LSA Count 3
Inter Area Pfx LSA Count 12000
Inter Area Rtr LSA Count 0
Group Mem LSA Count 0

```

```

Dell#show ipv6 ospf database grace-lsa
!
Type-11 Grace LSA (Area 0)

LS Age : 10
Link State ID : 6.16.192.66
Advertising Router : 100.1.1.1
LS Seq Number : 0x80000001
Checksum : 0x1DF1
Length : 36
Associated Interface : Gi 5/3
Restart Interval : 180
Restart Reason : Switch to Redundant Processor

```


## OSPFv3 Authentication Using IPsec

OSPFv3 authentication using IP security (IPsec) is supported the MXL switch.

Starting in Dell Networking OS version 8.4.2.0, OSPFv3 uses IPsec to provide authentication for OSPFv3 packets. IPsec authentication ensures security in the transmission of OSPFv3 packets between IPsec-enabled routers.

IPsec is a set of protocols developed by the internet engineering task force (IETF) to support secure exchange of packets at the IP layer. IPsec supports two encryption modes: transport and tunnel.

- **Transport mode** — encrypts only the data portion (payload) of each packet, but leaves the header untouched.
- **Tunnel mode** — is more secure and encrypts both the header and payload. On the receiving side, an IPsec-compliant device decrypts each packet.

 **NOTE:** The Dell Networking OS supports only Transport Encryption mode in OSPFv3 authentication with IPsec.

With IPsec-based authentication, Crypto images are used to include the IPsec secure socket application programming interface (API) required for use with OSPFv3.

To ensure integrity, data origin authentication, detection and rejection of replays, and confidentiality of the packet, RFC 4302 and RFC 4303 propose using two security protocols — authentication header (AH) and encapsulating security payload (ESP). For OSPFv3, these two IPsec protocols provide interoperable, high-quality cryptographically-based security.

- **HA** — IPsec authentication header is used in packet authentication to verify that data is not altered during transmission and ensures that users are communicating with the intended individual or organization. Insert the authentication header after the IP header with a value of 51. AH provides integrity and validation of data origin by authenticating every OSPFv3 packet. For detailed information about the IP AH protocol, refer to *RFC 4302*.
- **ESP** — encapsulating security payload encapsulates data, enabling the protection of data that follows in the datagram. ESP provides authentication and confidentiality of every packet. The ESP extension header is designed to provide a combination of security services for both IPv4 and IPv6. Insert the ESP header after the IP header and before the next layer protocol header in Transport mode. It is possible to insert the ESP header between the next layer protocol header and encapsulated IP header in Tunnel mode. However, Tunnel mode is not supported in the Dell Networking OS. For detailed information about the IP ESP protocol, refer to *RFC 4303*.

In OSPFv3 communication, IPsec provides security services between a pair of communicating hosts or security gateways using either AH or ESP. In an authentication policy on an interface or in an OSPF area, AH and ESP are used alone; in an encryption policy, AH and ESP may be used together. The difference between the two mechanisms is the extent of the coverage. ESP only protects IP header fields if they are encapsulated by ESP.


You decide the set of IPsec protocols that are employed for authentication and encryption and the ways in which they are employed. When you correctly implement and deploy IPsec, it does not adversely affect users or hosts. AH and ESP are designed to be cryptographic algorithm-independent.

### OSPFv3 Authentication Using IPsec: Configuration Notes

OSPFv3 authentication using IPsec is implemented according to the specifications in RFC 4552.

- To use IPsec, configure an authentication (using AH) or encryption (using ESP) security policy on an interface or in an OSPFv3 area. Each security policy consists of a security policy index (SPI) and the key used to validate OSPFv3 packets. After IPsec is configured for OSPFv3, IPsec operation is invisible to the user.
  - You can only enable one security protocol (AH or ESP) at a time on an interface or for an area. Enable IPsec AH with the `ipv6 ospf authentication` command; enable IPsec ESP with the `ipv6 ospf encryption` command.
  - The security policy configured for an area is inherited by default on all interfaces in the area.
  - The security policy configured on an interface overrides any area-level configured security for the area to which the interface is assigned.
  - The configured authentication or encryption policy is applied to all OSPFv3 packets transmitted on the interface or in the area. The IPsec security associations (SAs) are the same on inbound and outbound traffic on an OSPFv3 interface.
  - There is no maximum AH or ESP header length because the headers have fields with variable lengths.
- Manual key configuration is supported in an authentication or encryption policy (dynamic key configuration using the internet key exchange [IKE] protocol is not supported).
- In an OSPFv3 authentication policy:

- AH is used to authenticate OSPFv3 headers and certain fields in IPv6 headers and extension headers.
- MD5 and SHA1 authentication types are supported; encrypted and unencrypted keys are supported.
- In an OSPFv3 encryption policy:
  - Both encryption and authentication are used.
  - IPsec security associations (SAs) are supported only in Transport mode (Tunnel mode is not supported).
  - ESP with null encryption is supported for authenticating only OSPFv3 protocol headers.
  - ESP with non-null encryption is supported for full confidentiality.
  - 3DES, DES, AES-CBC, and NULL encryption algorithms are supported; encrypted and unencrypted keys are supported.

 **NOTE:** To encrypt all keys on a router, use the `service password-encryption` command in Global Configuration mode. However, this command does not provide a high level of network security. To enable key encryption in an IPsec security policy at an interface or area level, specify 7 for `[key-encryption-type]` when you enter the `ipv6 ospf authentication ipsec` or `ipv6 ospf encryption ipsec` command.

- To configure an IPsec security policy for authenticating or encrypting OSPFv3 packets on a physical, port-channel, or VLAN interface or OSPFv3 area, perform any of the following tasks:
  - [Configuring IPsec Authentication on an Interface](#)
  - [Configuring IPsec Encryption on an Interface](#)
  - [Configuring IPsec Authentication for an OSPFv3 Area](#)
  - [Configuring IPsec Encryption for an OSPFv3 Area](#)
  - [Displaying OSPFv3 IPsec Security Policies](#)

## Configuring IPsec Authentication on an Interface

To configure, remove, or display IPsec authentication on an interface, use the following commands.

**Prerequisite:** Before you enable IPsec authentication on an OSPFv3 interface, first enable IPv6 unicast routing globally, configure an IPv6 address and enable OSPFv3 on the interface, and assign it to an area (refer to [Configuration Task List for OSPFv2 \(OSPF for IPv4\)](#)).

The SPI value must be unique to one IPsec security policy (authentication or encryption) on the router. Configure the same authentication policy (the same SPI and key) on each OSPFv3 interface in a link.

- Enable IPsec authentication for OSPFv3 packets on an IPv6-based interface.

INTERFACE mode

```
ipv6 ospf authentication {null | ipsec spi number {MD5 | SHA1} [key-encryption-type] key}
```


- `null`: causes an authentication policy configured for the area to not be inherited on the interface.
- `ipsec spi number`: the security policy index (SPI) value. The range is from 256 to 4294967295.
- `MD5 | SHA1`: specifies the authentication type: Message Digest 5 (MD5) or Secure Hash Algorithm 1 (SHA-1).
- `key-encryption-type`: (optional) specifies if the key is encrypted. The valid values are 0 (key is not encrypted) or 7 (key is encrypted).

- `key`: specifies the text string used in authentication. All neighboring OSPFv3 routers must share key to exchange information. For MD5 authentication, the key must be 32 hex digits (non-encrypted) or 64 hex digits (encrypted). For SHA-1 authentication, the key must be 40 hex digits (non-encrypted) or 80 hex digits (encrypted).
- Remove an IPsec authentication policy from an interface.  
`no ipv6 ospf authentication ipsec spi number`
- Remove null authentication on an interface to allow the interface to inherit the authentication policy configured for the OSPFv3 area.  
`no ipv6 ospf authentication null`
- Display the configuration of IPsec authentication policies on the router.  
`show crypto ipsec policy`
- Display the security associations set up for OSPFv3 interfaces in authentication policies.  
`show crypto ipsec sa ipv6`

## Configuring IPsec Encryption on an Interface

To configure, remove, or display IPsec encryption on an interface, use the following commands.

**Prerequisite:** Before you enable IPsec encryption on an OSPFv3 interface, first enable IPv6 unicast routing globally, configure an IPv6 address and enable OSPFv3 on the interface, and assign it to an area (refer to [Configuration Task List for OSPFv2 \(OSPF for IPv4\)](#)).

 **NOTE:** When you configure encryption using the `ipv6 ospf encryption ipsec` command, you enable both IPsec encryption and authentication. However, when you enable authentication on an interface using the `ipv6 ospf authentication ipsec` command, you do not enable encryption at the same time.

The SPI value must be unique to one IPsec security policy (authentication or encryption) on the router. Configure the same authentication policy (the same SPI and key) on each OSPFv3 interface in a link.

- Enable IPsec encryption for OSPFv3 packets on an IPv6-based interface.

INTERFACE mode

```
ipv6 ospf encryption {null | ipsec spi number esp encryption-algorithm [key-encryption-type] key authentication-algorithm [key-authentication-type] key}
```

- `null`: causes an encryption policy configured for the area to not be inherited on the interface.
- `ipsec spi number`: is the security policy index (SPI) value. The range is from 256 to 4294967295.
- `esp encryption-algorithm`: specifies the encryption algorithm used with ESP. The valid values are 3DES, DES, AES-CBC, and NULL. For AES-CBC, only the AES-128 and AES-192 ciphers are supported.
- `key`: specifies the text string used in the encryption. All neighboring OSPFv3 routers must share the same key to decrypt information. Required lengths of a non-encrypted or encrypted key are: 3DES - 48 or 96 hex digits; DES - 16 or 32 hex digits; AES-CBC - 32 or 64 hex digits for AES-128 and 48 or 96 hex digits for AES-192.
- `key-encryption-type`: (optional) specifies if the key is encrypted. The valid values are 0 (key is not encrypted) or 7 (key is encrypted).
- `authentication-algorithm`: specifies the encryption authentication algorithm to use. The valid values are MD5 or SHA1.
- `key`: specifies the text string used in authentication. All neighboring OSPFv3 routers must share key to exchange information. For MD5 authentication, the key must be 32 hex digits (non-encrypted) or 64 hex digits (encrypted). For SHA-1 authentication, the key must be 40 hex digits (non-encrypted) or 80 hex digits (encrypted).

- *key-authentication-type*: (optional) specifies if the authentication key is encrypted. The valid values are 0 or 7.
- Remove an IPsec encryption policy from an interface.
 

```
no ipv6 ospf encryption ipsec spi number
```
- Remove null encryption on an interface to allow the interface to inherit the encryption policy configured for the OSPFv3 area.
 

```
no ipv6 ospf encryption null
```
- Display the configuration of IPsec encryption policies on the router.
 

```
show crypto ipsec policy
```
- Display the security associations set up for OSPFv3 interfaces in encryption policies.
 

```
show crypto ipsec sa ipv6
```

## Configuring IPsec Authentication for an OSPFv3 Area

To configure, remove, or display IPsec authentication for an OSPFv3 area, use the following commands.

**Prerequisite:** Before you enable IPsec authentication on an OSPFv3 area, first enable OSPFv3 globally on the router (refer to [Configuration Task List for OSPFv2 \(OSPF for IPv4\)](#)).

The security policy index (SPI) value must be unique to one IPsec security policy (authentication or encryption) on the router. Configure the same authentication policy (the same SPI and key) on each interface in an OSPFv3 link.

If you have enabled IPsec encryption in an OSPFv3 area using the `area encryption` command, you cannot use the `area authentication` command in the area at the same time.

The configuration of IPsec authentication on an interface-level takes precedence over an area-level configuration. If you remove an interface configuration, an area authentication policy that has been configured is applied to the interface.

- Enable IPsec authentication for OSPFv3 packets in an area.

CONF-IPV6-ROUTER-OSPF mode

```
area-id authentication ipsec spi number {MD5 | SHA1} [key-encryption-type]
key
```

- *area area-id*: specifies the area for which OSPFv3 traffic is to be authenticated. For *area-id*, enter a number or an IPv6 prefix.
- *spi number*: is the SPI value. The range is from 256 to 4294967295.
- MD5 | SHA1: specifies the authentication type: message digest 5 (MD5) or Secure Hash Algorithm 1 (SHA-1).
- *key-encryption-type*: (optional) specifies if the key is encrypted. The valid values are 0 (key is not encrypted) or 7 (key is encrypted).
- *key*: specifies the text string used in authentication. All neighboring OSPFv3 routers must share key to exchange information. For MD5 authentication, the key must be 32 hex digits (non-encrypted) or 64 hex digits (encrypted). For SHA-1 authentication, the key must be 40 hex digits (non-encrypted) or 80 hex digits (encrypted).
- Remove an IPsec authentication policy from an OSPFv3 area.
 

```
no area area-id authentication ipsec spi number
```
- Display the configuration of IPsec authentication policies on the router.
 

```
show crypto ipsec policy
```




## Configuring IPsec Encryption for an OSPFv3 Area

To configure, remove, or display IPsec encryption in an OSPFv3 area, use the following commands.

**Prerequisite:** Before you enable IPsec encryption in an OSPFv3 area, first enable OSPFv3 globally on the router (refer to [Configuration Task List for OSPFv3 \(OSPF for IPv6\)](#)).

The SPI value must be unique to one IPsec security policy (authentication or encryption) on the router.

Configure the same encryption policy (the same SPI and keys) on each interface in an OPSFv3 link.

 **NOTE:** When you configure encryption using the `area encryption` command, you enable both IPsec encryption and authentication. However, when you enable authentication on an area using the `area authentication` command, you do not enable encryption at the same time.

If you have enabled IPsec authentication in an OSPFv3 area using the `area authentication` command, you cannot use the `area encryption` command in the area at the same time.

The configuration of IPsec encryption on an interface-level takes precedence over an area-level configuration. If you remove an interface configuration, an area encryption policy that has been configured is applied to the interface.

- Enable IPsec encryption for OSPFv3 packets in an area.

CONF-IPV6-ROUTER-OSPF mode

```
area area-id encryption ipsec spi number esp encryption-algorithm [key-
encryption-type] key authentication-algorithm [key-authentication-type] key
```

- `area area-id`: specifies the area for which OSPFv3 traffic is to be encrypted. For `area-id`, enter a number or an IPv6 prefix.
- `spi number`: is the security policy index (SPI) value. The range is from 256 to 4294967295.
- `esp encryption-algorithm`: specifies the encryption algorithm used with ESP. The valid values are 3DES, DES, AES-CBC, and NULL. For AES-CBC, only the AES-128 and AES-192 ciphers are supported.
- `key`: specifies the text string used in the encryption. All neighboring OSPFv3 routers must share the same key to decrypt information. The required lengths of a non-encrypted or encrypted key are: 3DES - 48 or 96 hex digits; DES - 16 or 32 hex digits; AES-CBC - 32 or 64 hex digits for AES-128 and 48 or 96 hex digits for AES-192.
- `key-encryption-type`: (optional) specifies if the key is encrypted. Valid values: 0 (key is not encrypted) or 7 (key is encrypted).
- `authentication-algorithm`: specifies the authentication algorithm to use for encryption. The valid values are MD5 or SHA1.
- `key`: specifies the text string used in authentication. All neighboring OSPFv3 routers must share key to exchange information. For MD5 authentication, the key must be 32 hex digits (non-encrypted) or 64 hex digits (encrypted). For SHA-1 authentication, the key must be 40 hex digits (non-encrypted) or 80 hex digits (encrypted).
- `key-authentication-type`: (optional) specifies if the authentication key is encrypted. The valid values are 0 or 7.

- Remove an IPsec encryption policy from an OSPFv3 area.

```
no area area-id encryption ipsec spi number
```

- Display the configuration of IPsec encryption policies on the router.

```
show crypto ipsec policy
```

## Displaying OSPFv3 IPsec Security Policies

To display the configuration of IPsec authentication and encryption policies, use the following commands.

- Display the AH and ESP parameters configured in IPsec security policies, including the SPI number, key, and algorithms used.

EXEC Privilege mode

```
show crypto ipsec policy [name name]
```

– *name*: displays configuration details about a specified policy.

- Display security associations set up for OSPFv3 links in IPsec authentication and encryption policies on the router.

EXEC Privilege

```
show crypto ipsec sa ipv6 [interface interface]
```

To display information on the SAs used on a specific interface, enter `interface interface`, where *interface* is one of the following values:

- For a Port Channel interface, enter the keywords `port-channel number`.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a 40-Gigabit Ethernet interface, enter the keyword `FortyGigabitEthernet` then the slot/port information.
- For a VLAN interface, enter the keywords `vlan vlan-id`. The valid VLAN IDs are from 1 to 4094.

### Example of the `show crypto ipsec policy` Command

### Example of the `show crypto ipsec sa ipv6` Command

In the first example, the keys are not encrypted (shown in bold). In the second and third examples, the keys are encrypted (shown in bold).

```
Dell#show crypto ipsec policy
```

```
Crypto IPsec client security policy data
```

```
Policy name : OSPFv3-1-502
Policy refcount : 1
Inbound ESP SPI : 502 (0x1F6)
Outbound ESP SPI : 502 (0x1F6)
Inbound ESP Auth Key : 123456789a123456789b123456789c12
Outbound ESP Auth Key : 123456789a123456789b123456789c12
Inbound ESP Cipher Key : 123456789a123456789b123456789c123456789d12345678
Outbound ESP Cipher Key : 123456789a123456789b123456789c123456789d12345678
Transform set : esp-3des esp-md5-hmac
```

```
Crypto IPsec client security policy data
```

```
Policy name : OSPFv3-1-500
Policy refcount : 2
Inbound AH SPI : 500 (0x1F4)
Outbound AH SPI : 500 (0x1F4)
Inbound AH Key :
bbdd96e6eb4828e2e27bc3f9ff541e43faa759c9ef5706ba8ed8bb5efe91e97e
Outbound AH Key :
```

```
bbdd96e6eb4828e2e27bc3f9ff541e43faa759c9ef5706ba8ed8bb5efe91e97e
Transform set : ah-md5-hmac
```

Crypto IPsec client security policy data

```
Policy name : OSPFv3-0-501
Policy refcount : 1
Inbound ESP SPI : 501 (0x1F5)
Outbound ESP SPI : 501 (0x1F5)
Inbound ESP Auth Key :
bbdd96e6eb4828e2e27bc3f9ff541e43faa759c9ef5706ba8ed8bb5efe91e97eb7c0c30808825fb5
Outbound ESP Auth Key :
bbdd96e6eb4828e2e27bc3f9ff541e43faa759c9ef5706ba8ed8bb5efe91e97eb7c0c30808825fb5
Inbound ESP Cipher Key :
bbdd96e6eb4828e2e27bc3f9ff541e43faa759c9ef5706ba10345a1039ba8f8a
Outbound ESP Cipher Key :
bbdd96e6eb4828e2e27bc3f9ff541e43faa759c9ef5706ba10345a1039ba8f8a
Transform set : esp-128-aes esp-sha1-hmac
```

Dell#show crypto ipsec sa ipv6

```
Interface: TenGigabitEthernet 0/0
Link Local address: fe80::201:e8ff:fe40:4d10
IPSecv6 policy name: OSPFv3-1-500
```

```
inbound ah sas
spi : 500 (0x1f4)
transform : ah-md5-hmac
in use settings : {Transport, }
replay detection support : N
STATUS : ACTIVE
```

```
outbound ah sas
spi : 500 (0x1f4)
transform : ah-md5-hmac
in use settings : {Transport, }
replay detection support : N
STATUS : ACTIVE
```

```
inbound esp sas
```

```
outbound esp sas
```

```
Interface: TenGigabitEthernet 0/1
Link Local address: fe80::201:e8ff:fe40:4d11
IPSecv6 policy name: OSPFv3-1-600
```

```
inbound ah sas
```

```
outbound ah sas
```


```
inbound esp sas
spi : 600 (0x258)
transform : esp-des esp-sha1-hmac
in use settings : {Transport, }
replay detection support : N
STATUS : ACTIVE
```

```
outbound esp sas
spi : 600 (0x258)
transform : esp-des esp-sha1-hmac
in use settings : {Transport, }
```

```
replay detection support : N
STATUS : ACTIVE
```

## Troubleshooting OSPFv3

The Dell Networking OS has several tools to make troubleshooting easier. Consider the following information as these are typical issues that interrupt the OSPFv3 process.

 **NOTE:** The following troubleshooting section is not meant to be a comprehensive list, only examples of typical troubleshooting checks.

- Have you enabled OSPF globally?
- Is the OSPF process active on the interface?
- Are the adjacencies established correctly?
- Did you configure the interfaces for Layer 3 correctly?
- Is the router in the correct area type?
- Did you include the routes in the OSPF database?
- Did you include the OSPF routes in the routing table (not just the OSPF database)?

Some useful troubleshooting commands are:

- `show ipv6 interfaces`
- `show ipv6 protocols`
- `debug ipv6 ospf events and/or packets`
- `show ipv6 neighbors`
- `show virtual links`
- `show ipv6 routes`

## Viewing Summary Information

To get general route, configuration, links status, and debug information, use the following commands.

- View the summary information of the IPv6 routes.  
EXEC Privilege mode  
  
`show ipv6 route summary`
- View the summary information for the OSPFv3 database.  
EXEC Privilege mode  
  
`show ipv6 ospf database`
- View the configuration of OSPFv3 neighbors.  
EXEC Privilege mode  
  
`show ipv6 ospf neighbor`
- View debug messages for all OSPFv3 interfaces.  
EXEC Privilege mode  
  
`debug ipv6 ospf [event | packet] {type slot/port}`  
  
– event: View OSPF event messages.

- `packet`: View OSPF packets.
- For a Gigabit Ethernet interface, enter the keyword `GigabitEthernet` then the slot/port information (for example, `passive-interface gi 2/1`).
- For a port channel, enter the keywords `port-channel` then a number from 1 to 255.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information (for example, `passive-interface ten 2/3`).
- For a 40-Gigabit Ethernet interface, enter the keyword `FortyGigabitEthernet` then the slot/port information.
- For a VLAN, enter the keyword `vlan` then a number from 1 to 4094 (for example, `passive-interface vlan 2222`). The system supports 4094 VLANs.

# Policy-based Routing (PBR)

Policy-based Routing is supported on the MXL platform.

This chapter covers the following topics:

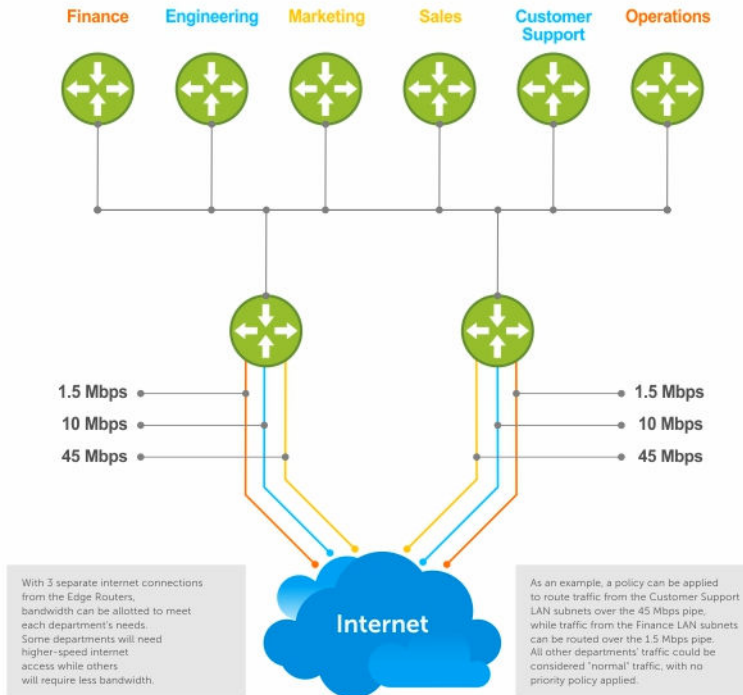
- Overview
- Implementing Policy-based Routing with Dell Networking OS
- Configuration Task List for Policy-based Routing
- Sample Configuration

## Overview

Policy-based Routing (PBR) enables you to make routing decisions based on policies applied to a specific interface. When a router receives a packet it normally decides where to forward it based on the destination address in the packet, which is used to look up an entry in a routing table. However, in some cases, there may be a need to forward the packet based on other criteria: size, source, protocol type, destination, etc. For example, a network administrator might want to forward a packet that uses TCP across a different next-hop than packets using ICMP.

Rules for **PBR** can also be a combination of things:

When the packet comes from this source and wants to go to that destination then route it to next-hop or onto that specific interface. This permits routing over different links or towards different networks even while the destination is the same but depending on where the packet originates.



To enable a PBR, you create a Redirect List. Redirect lists are defined by rules, or routing policies. The following parameters can be defined in the routing policies or rules:

- IP address of the forwarding router (next-hop IP address)
- Protocol as defined in the header
- Source IP address and mask
- Destination IP address and mask
- Source port
- Destination port
- TCP Flags

Once a redirect-list is applied to an interface, all traffic passing through it is subjected to the rules defined in the redirect-list.

The traffic is forwarded based on the following:

- Next-hop addresses are verified. If the specified next hop is reachable, then the traffic is forwarded to the specified next-hop.
- If the specified next-hops are not reachable, then the normal routing table is used to forward the traffic.
- Dell Networking OS supports multiple next-hop entries in the redirect lists.
- Redirect-Lists are applied at Ingress.

PBR with Redirect-to-Tunnel Option:

The user can provide a tunnel id for a redirect rule. In this case, the resolved next hop would be the tunnel interface IP. The qualifiers of the rule would be pertaining to the inner IP details. For next hop to be

a tunnel interface user needs to provide tunnel id mandatory. Instead if user provides the tunnel destination IP as next hop, that would be treated as IPv4 next hop and not tunnel next hop.

PBR with Multiple Tacking Option:

Policy based routing with multiple tracking option extends and introduces the capabilities of object tracking to verify the next hop IP address before forwarding the traffic to the next hop. The verification method is made transparent to the user. The multiple tracking options feature is most suitable for routers which have multiple devices as the next hop (primarily indirect next-hops and/or Tunnel Interfaces in this case). It allows you to backup Indirect Next-hop with another, choose the specific Indirect Next-hop and/or Tunnel Interface which is available by sending ICMP pings to verify reach ability and/or check the Tunnel Interface UP or DOWN status, and then route traffic out to the next-hop and/or Tunnel Interface.

## Implementing Policy-based Routing with Dell Networking OS

- Non-contiguous bitmasks for PBR
- Hot-Lock PBR

### Non-contiguous bitmasks for PBR

Non-contiguous bitmasks for PBR allows more granular and flexible control over routing policies. Network addresses that are in the middle of a subnet can be included or excluded. Specific bitmasks can be entered using the dotted decimal format.

### Non-contiguous bitmask example

```
Dell#show ip redirect-list
IP redirect-list rcl0:
Defined as:
seq 5 permit ip 200.200.200.200 200.200.200.200 199.199.199.199 199.199.199.199
seq 10 redirect 1.1.1.2 tcp 234.224.234.234 255.234.234.234 222.222.222.222/24
seq 40 ack, Next-hop reachable(via Te 8/1)
Applied interfaces:
Te 8/0
```

### Hot-Lock PBR

Ingress and egress Hot Lock PBR allow you to add or delete new rules into an existing policy (already written into CAM) without disruption to traffic flow. Existing entries in CAM are adjusted to accommodate the new entries. Hot Lock PBR is enabled by default.

## Configuration Task List for Policy-based Routing

To enable the PBR:

- Create a Redirect List
- Create a Rule for a Redirect-list
- Create a Track-id list. For complete tracking information, refer to Object Tracking chapter.
- Apply a Redirect-list to an Interface using a Redirect-group



## Create a Redirect List

Use the following command in **CONFIGURATION** mode:

| Command Syntax                                    | Command Mode  | Purpose                                                                    |
|---------------------------------------------------|---------------|----------------------------------------------------------------------------|
| <b>ip redirect-list</b> <i>redirect-list-name</i> | CONFIGURATION | Create a redirect list by entering the list name.<br>Format: 16 characters |

Delete the redirect list with the **no ip redirect-list** command.

The following example creates a redirect list by the name of "xyz."

```
Dell(conf)#ip redirect-list ?
WORD Redirect-list name (max 16 chars)
Dell(conf)#ip redirect-list xyz
```

## Create a Rule for a Redirect-list

Use the following command in **CONFIGURATION REDIRECT-LIST** mode to set the rules for the redirect list. You can enter the command multiple times and create a sequence of redirect rules. Use the **seq nn redirect** version of the command to organize your rules.

| Command Syntax                                                                                                                                                                                                                                                                    | Command Mode       | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>seq</b> { <i>number</i> } <b>redirect</b> { <i>ip-address</i> }{ <i>ip-protocol-number</i>   <i>protocol-type</i> [ <i>bit</i> ]} { <i>source mask</i>   <b>any</b>   <b>host</b> <i>ip-address</i> } { <i>destination mask</i>   <b>any</b>   <b>host</b> <i>ip-address</i> } | CONF-REDIRECT-LIST | Configure a rule for the redirect list.<br><br><i>number</i> is the number in sequence to initiate this rule<br><br><i>ip-address</i> is the Forwarding router's address<br><br>FORMAT: A.B.C.D<br><br>FORMAT: slot/port<br><br><i>ip-protocol-number</i> or <i>protocol-type</i> is the type of protocol to be redirected<br><br>FORMAT: 0-255 for IP protocol number, or enter protocol type<br><br><i>source ip-address</i> or <i>any</i> or <i>host ip-address</i> is the Source's IP address<br><br>FORMAT: A.B.C.D/NN, or ANY or HOST IP address<br><br><i>destination ip-address</i> or <i>any</i> or <i>host ip-address</i> is the Destination's IP address<br><br>FORMAT: A.B.C.D/NN, or ANY or HOST IP address |

Delete a rule with the **no redirect** command.

The redirect rule supports Non-contiguous bitmasks for PBR in the Destination router IP address

The below step shows a step-by-step example of how to create a rule for a redirect list by configuring:

- IP address of the next-hop router in the forwarding route
- IP protocol number
- Source address with mask information
- Destination address with mask information

#### Creating a Rule Example:

```
Dell(conf-redirect-list)#redirect ?
A.B.C.D Forwarding router's address

Dell(conf-redirect-list)#redirect 3.3.3.3 ?
<0-255> An IP protocol number
icmp Internet Control Message Protocol
ip Any Internet Protocol
tcp Transmission Control Protocol
udp User Datagram Protocol
Dell(conf-redirect-list)#redirect 3.3.3.3 ip ?
A.B.C.D Source address
any Any source host
host A single source host
Dell(conf-redirect-list)#redirect 3.3.3.3 ip 222.1.1.1 ?
Mask A.B.C.D or /nn Mask in dotted decimal or in slash format
Dell(conf-redirect-list)#redirect 3.3.3.3 ip 222.1.1.1 /32 ?
A.B.C.D Destination address
any Any destination host
host A single destination host
Dell(conf-redirect-list)#redirect 3.3.3.3 ip 222.1.1.1 /32 77.1.1.1 ?
Mask A.B.C.D or /nn Mask in dotted decimal or in slash format
Dell(conf-redirect-list)#redirect 3.3.3.3 ip 222.1.1.1 /32 77.1.1.1 /32 ?
Dell(conf-redirect-list)#redirect 3.3.3.3 ip 222.1.1.1 /32 77.1.1.1 /32
Dell(conf-redirect-list)#do show ip redirect-list
```


```
IP redirect-list xyz:
 Defined as:
 seq 5 redirect 3.3.3.3 ip host 222.1.1.1 host 77.1.1.1
 Applied interfaces:
 None
```

Multiple rules can be applied to a single redirect-list. The rules are applied in ascending order, starting with the rule that has the lowest sequence number in a redirect-list displays the correct method for applying multiple rules to one list.

#### Creating multiple rules for a redirect-list:

```
Dell(conf)#ip redirect-list test
Dell(conf-redirect-list)#seq 10 redirect 10.1.1.2 ip 20.1.1.0/24 any
Dell(conf-redirect-list)#seq 15 redirect 10.1.1.3 ip 20.1.1.0/25 any
Dell(conf-redirect-list)#seq 20 redirect 10.1.1.3 ip 20.1.1.128/24 any
Dell(conf-redirect-list)#show config
!
ip redirect-list test
 seq 10 redirect 10.1.1.2 ip 20.1.1.0/24 any
 seq 15 redirect 10.1.1.3 ip 20.1.1.0/25 any
```

```
seq 20 redirect 10.1.1.3 ip 20.1.1.0/24 any
Dell(conf-redirect-list)#
```

 **NOTE:** Starting in release 9.4(0.0), Dell Networking OS supports the use of multiple recursive routes with the same source-address and destination-address combination in a redirect policy on an router.

A recursive route is a route for which the immediate next-hop address is learned dynamically through a routing protocol and acquired through a route lookup in the routing table. The user can configure multiple recursive routes in a redirect list by entering multiple **seq redirect** statements with the same source and destination address and specify a different next-hop IP address. In this way, the recursive routes are used as different forwarding routes for dynamic failover. If the primary path goes down and the recursive route is removed from the routing table, the **seq redirect** statement is ignored and the next statement in the list with a different route is used.

## PBR Exceptions (Permit)

Use the command **permit** to create an exception to a redirect list. Exceptions are used when a forwarding decision should be based on the routing table rather than a routing policy.

Dell Networking OS assigns the first available sequence number to a rule configured without a sequence number and inserts the rule into the PBR CAM region next to the existing entries. Since the order of rules is important, ensure that you configure any necessary sequence numbers.

The permit statement is never applied because the redirect list covers all source and destination IP addresses.

### Ineffective PBR Exception due to Low Sequence Number


```
ip redirect-list rcl0
seq 5 redirect 2.2.2.2 ip any any
seq 10 permit ip host 3.3.3.3 any
```

To ensure that the permit statement or PBR exception is effective, use a lower sequence number, as shown below:

```
ip redirect-list rcl0
seq 10 permit ip host 3.3.3.3 any
seq 15 redirect 2.2.2.2 ip any any
```

### Apply a Redirect-list to an Interface using a Redirect-group

IP redirect lists are supported on physical interfaces as well as VLAN and port-channel interfaces.

 **NOTE:** When you apply a redirect-list on a port-channel, when traffic is redirected to the next hop and the destination port-channel is shut down, the traffic is dropped. However, on the S-Series, the traffic redirected to the destination port-channel is sometimes switched.

Use the following command in **INTERFACE** mode to apply a redirect list to an interface. Multiple redirect-lists can be applied to a redirect-group. It is also possible to create two or more redirect-groups on one interface for backup purposes.

| Command Syntax                                     | Command Mode | Purpose                                                                                                                                                                                      |
|----------------------------------------------------|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>ip redirect-group</b> <i>redirect-list-name</i> | INTERFACE    | Apply a redirect list (policy-based routing) to an interface.<br><br><i>redirect-list-name</i> is the name of a redirect list to apply to this interface.<br><br>FORMAT: up to 16 characters |
|                                                    |              | Delete the redirect list from this interface with the <b>[no] ip redirect-group</b> command.                                                                                                 |

In this example, the list "xyz" is applied to the tenGigabitEthernet 4/0 interface.

#### Applying a Redirect-list to an Interface Example:

```
Dell(conf-if-te-4/0)#ip redirect-group xyz
Dell(conf-if-te-4/0)#
```

#### Applying a Redirect-list to an Interface Example:

```
Dell(conf-if-te-1/0)#ip redirect-group test
Dell(conf-if-te-1/0)#ip redirect-group xyz
Dell(conf-if-te-1/0)#show config
!
interface TenGigabitEthernet 1/0
 no ip address
 ip redirect-group test
 ip redirect-group xyz
 shutdown
Dell(conf-if-te-1/0)#
```

In addition to supporting multiple redirect-lists in a redirect-group, multiple redirect-groups are supported on a single interface. Dell Networking OS has the capability to support multiple groups on an interface for backup purposes.

#### Show Redirect List Configuration

To view the configuration redirect list configuration, use the following command in EXEC mode:

| Command Syntax                                         | Command Mode | Purpose                                                             |
|--------------------------------------------------------|--------------|---------------------------------------------------------------------|
| <b>show ip redirect-list</b> <i>redirect-list-name</i> | EXEC         | View the redirect list configuration and the associated interfaces. |
| <b>show cam pbr</b>                                    | EXEC         | View the redirect list entries programmed in the CAM.               |
| <b>show cam-usage</b>                                  |              |                                                                     |

List the redirect list configuration using the **show ip redirect-list redirect-list-name** command. The non-contiguous mask is displayed in dotted format (x.x.x.x). The contiguous mask is displayed in /x format. Some sample outputs are shown below:

```
Dell#show ip redirect-list explicit_tunnel
IP redirect-list explicit_tunnel:
```

```

Defined as:
seq 5 redirect tunnel 1 track 1 tcp 155.55.2.0/24 222.22.2.0/24, Track 1 [up],
Next-hop reachable (via Te 1/32)
seq 10 redirect tunnel 1 track 1 tcp any any, Track 1 [up], Next-hop reachable
(via Te 1/32)
seq 15 redirect tunnel 2 udp 155.55.0.0/16 host 144.144.144.144, Track 1 [up],
Next-hop reachable (via Te 1/32)
seq 35 redirect 155.1.1.2 track 5 ip 7.7.7.0/24 8.8.8.0/24, Track 5 [up], Next-
hop reachable (via Po 5)
seq 30 redirect 155.1.1.2 track 6 icmp host 8.8.8.8 any, Track 5 [up], Next-hop
reachable (via Po 5)
seq 35 redirect 42.1.1.2 icmp host 8.8.8.8 any, Next-hop reachable (via Vl 20)
seq 40 redirect 43.1.1.2 tcp 155.55.2.0/24 222.22.2.0/24, Next-hop reachable
(via Vl 30)
seq 45 redirect 31.1.1.2 track 200 ip 12.0.0.0 255.0.0.197 13.0.0.0
255.0.0.197, Track 200 [up], Next-hop reachable (via Te 1/32)

Track 200 [up], Next-hop reachable (via Vl 20)
Track 200 [up], Next-hop reachable (via Po 5)
Track 200 [up], Next-hop reachable (via Po 7)
Track 200 [up], Next-hop reachable (via Te 2/18)
Track 200 [up], Next-hop reachable (via Te 2/19)

```


Use the **show ip redirect-list** (without the list name) to display all the redirect-lists configured on the device.

```
Dell#show ip redirect-list
```

```

IP redirect-list rcl0:
Defined as:
seq 5 permit ip 200.200.200.200 200.200.200.200 199.199.199.199 199.199.199.199
seq 10 redirect 1.1.1.2 tcp 234.224.234.234 255.234.234.234 222.222.222.222/24
seq 40 ack, Next-hop reachable (via Te 8/1), ARP resolved
Applied interfaces:
Te 8/0

```

 **NOTE:** If, the redirect-list is applied to an interface, the output of **show ip redirect-list redirect-list-name** command displays reachability and ARP status for the specified next-hop.

#### Showing CAM PBR Configuration Example :

```
Dell#show cam pbr stack-unit 1 port-set 0
```

```
TCP Flag: Bit 5 - URG, Bit 4 - ACK, Bit 3 - PSH, Bit 2 - RST, Bit 1 - SYN, Bit
0 - FIN
```

| Cam Egress Index | Port Flag | VlanID Port | Proto Port | Tcp MAC | Src Port | Dst | SrcIp           | DstIp                  | Next-hop                      |
|------------------|-----------|-------------|------------|---------|----------|-----|-----------------|------------------------|-------------------------------|
| 06080            | 0         | N/A         | IP         | 0x0     | 0        | 0   | 200.200.200.200 | <b>200.200.200.200</b> | 199.199.199.199               |
|                  |           |             |            |         |          |     |                 |                        | <b>199.199.199.199</b> N/A NA |
| 06081            | 0         | N/A         | TCP        | 0x10    | 0        | 40  | 234.234.234.234 | 255.234.234.234        | 222.222.222.222/24            |
|                  |           |             |            |         |          |     |                 |                        | 00:00:00:00:00:09 8/1         |

## Sample Configuration


The following configuration is an example for setting up a PBR. These are not comprehensive directions. They are intended to give you a some guidance with typical configurations. You can copy and paste from these examples to your CLI. Be sure you make the necessary changes to support your own IP Addresses, Interfaces, Names, etc.

Graphic illustration of the configuration shown below:

The Redirect-List GOLD defined in this example, creates the following rules:

- description Route Gold traffic to the DS3.
- seq 5 redirect 10.99.99.254 ip 192.168.1.0/24 any " Redirect to next-hop router IP 10.99.99.254 any traffic originating in 192.168.1.0/24"
- seq 10 redirect 10.99.99.254 ip 192.168.2.0/24 any " Redirect to next-hop router IP 10.99.99.254 any traffic originating in 192.168.2.0/24"
- seq 15 permit ip any

PBR Sample Configuration examples are shown below:



### Create the Redirect-List GOLD

```
EDGE_ROUTER(conf-if-Te-3/23)#ip redirect-list GOLD
EDGE_ROUTER(conf-redirect-list)#description Route GOLD traffic to ISP_GOLD.
```

```

EDGE_ROUTER(conf-redirect-list)#direct 10.99.99.254 ip 192.168.1.0/24 any
EDGE_ROUTER(conf-redirect-list)#redirect 10.99.99.254 ip 192.168.2.0/24 any
EDGE_ROUTER(conf-redirect-list)# seq 15 permit ip any any
EDGE_ROUTER(conf-redirect-list)#show config
!
ip redirect-list GOLD
description Route GOLD traffic to ISP GOLD.
seq 5 redirect 10.99.99.254 ip 192.168.1.0/24 any
seq 10 redirect 10.99.99.254 ip 192.168.2.0/24 any
seq 15 permit ip any any

```

## Assign Redirect-List GOLD to Interface 2/11

```

EDGE_ROUTER(conf)#int Te 2/11
EDGE_ROUTER(conf-if-Te-2/11)#ip add 192.168.3.2/24
EDGE_ROUTER(conf-if-Te-2/11)#no shut
EDGE_ROUTER(conf-if-Te-2/11)#
EDGE_ROUTER(conf-if-Te-2/11)#ip redirect-group GOLD
EDGE_ROUTER(conf-if-Te-2/11)#no shut
EDGE_ROUTER(conf-if-Te-2/11)#end
EDGE_ROUTER(conf-redirect-list)#end

```

```
EDGE_ROUTER#
```

## View Redirect-List GOLD

```
EDGE_ROUTER#show ip redirect-list
```

```

IP redirect-list GOLD:
 Defined as:
 seq 5 redirect 10.99.99.254 ip 192.168.1.0/24 any, Next-hop reachable (via Te
3/23), ARP resolved
 seq 10 redirect 10.99.99.254 ip 192.168.2.0/24 any, Next-hop reachable (via
Te 3/23), ARP resolved
 seq 15 permit ip any any
 Applied interfaces:
 Te 2/11
EDGE_ROUTER#

```

## Configuration Tasks for Creating a PBR list using Explicit Track Objects for Redirect IP's

Create Track Objects to track the Redirect IP's:

```

Dell#configure terminal
Dell(conf)#track 3 ip host 42.1.1.2 reachability
Dell(conf-track-3)#probe icmp
Dell(conf-track-3)#track 4 ip host 43.1.1.2 reachability
Dell(conf-track-4)#probe icmp
Dell(conf-track-4)#end

```

Create a Redirect-list with Track Objects pertaining to Redirect-IP's:

```

Dell#configure terminal
Dell(conf)#ip redirect-list redirect_list_with_track
Dell(conf-redirect-list)#redirect 42.1.1.2 track 3 tcp 155.55.2.0/24
222.22.2.0/24
Dell(conf-redirect-list)#redirect 42.1.1.2 track 3 tcp any any
Dell(conf-redirect-list)#redirect 42.1.1.2 track 3 udp 155.55.0.0/16 host
144.144.144.144
Dell(conf-redirect-list)#redirect 42.1.1.2 track 3 udp any host 144.144.144.144

```

```
Dell(conf-redirect-list)#redirect 43.1.1.2 track 4 ip host 7.7.7.7 host
144.144.144.144
Dell(conf-redirect-list)#end
```

Verify the Status of the Track Objects (Up/Down):

```
Dell#show track brief
```

| ResId | Resource               | Parameter   | State | LastChange |
|-------|------------------------|-------------|-------|------------|
| 1     | Interface ip routing   | Tunnel 1    | Up    | 00:02:16   |
| 2     | Interface ipv6 routing | Tunnel 2    | Up    | 00:03:31   |
| 3     | IP Host reachability   | 42.1.1.2/32 | Up    | 00:00:59   |
| 4     | IP Host reachability   | 43.1.1.2/32 | Up    | 00:00:59   |

Apply the Redirect Rule to an Interface:

```
Dell#
Dell(conf)#int TenGigabitEthernet 2/28
Dell(conf-if-te-2/28)#ip redirect-group redirect_list_with_track
Dell(conf-if-te-2/28)#end
```

Verify the Applied Redirect Rules:

```
Dell#show ip redirect-list redirect_list_with_track

IP redirect-list redirect_list_with_track
Defined as:
 seq 5 redirect 42.1.1.2 track 3 tcp 155.55.2.0/24 222.22.2.0/24, Track 3
[up], Next-hop reachable (via V1 20)
 seq 10 redirect 42.1.1.2 track 3 tcp any any, Track 3 [up], Next-hop
reachable (via V1 20)
 seq 15 redirect 42.1.1.2 track 3 udp 155.55.0.0/16 host 144.144.144.144,
Track 3 [up], Next-hop reachable (via V1 20)
 seq 20 redirect 42.1.1.2 track 3 udp any host 144.144.144.144, Track 3 [up],
Next-hop reachable (via V1 20)
 seq 25 redirect 43.1.1.2 track 4 ip host 7.7.7.7 host 144.144.144.144, Track
4 [up], Next-hop reachable (via V1 20)
Applied interfaces:
 Te 2/28
Dell#
```

## Configuration Tasks for Creating a PBR list using Explicit Track Objects for Tunnel Interfaces

Creating steps for Tunnel Interfaces:

```
Dell#configure terminal
Dell(conf)#interface tunnel 1
Dell(conf-if-tu-1)#tunnel destination 40.1.1.2
Dell(conf-if-tu-1)#tunnel source 40.1.1.1
Dell(conf-if-tu-1)#tunnel mode ipip
Dell(conf-if-tu-1)#tunnel keepalive 60.1.1.2
Dell(conf-if-tu-1)#ip address 60.1.1.1/24
Dell(conf-if-tu-1)#ipv6 address 600:10::1/64
Dell(conf-if-tu-1)#no shutdown
Dell(conf-if-tu-1)#end
Dell#
```

```
Dell#configure terminal
Dell(conf)#interface tunnel 2
Dell(conf-if-tu-2)#tunnel destination 441:10::2
Dell(conf-if-tu-2)#tunnel source 441:10::1
Dell(conf-if-tu-2)#tunnel mode ipv6
```



```
Dell(conf-if-tu-2)#tunnel keepalive 601:10::2
Dell(conf-if-tu-2)#ipv6 address 601:10::1/64
Dell(conf-if-tu-2)#no shutdown
Dell(conf-if-tu-2)#end
Dell#
```

Create Track Objects to track the Tunnel Interfaces:

```
Dell#configure terminal
Dell(conf)#track 1 interface tunnel 1 ip routing
Dell(conf-track-1)#exit
Dell(conf)#track 2 interface tunnel 2 ipv6 routing
Dell(conf-track-2)#end
```

Verify the Status of the Track Objects (Up/Down):

```
Dell#show track brief
```

```
ResId Resource Parameter State LastChange
1 Interface ip routing Tunnel 1 Up 00:00:00
2 Interface ipv6 routing Tunnel 2 Up 00:00:00
Dell#
```

Create a Redirect-list with Track Objects pertaining to Tunnel Interfaces:

```
Dell#configure terminal
Dell(conf)#ip redirect-list explicit_tunnel
Dell(conf-redirect-list)#redirect tunnel 1 track 1 tcp 155.55.2.0/24
222.22.2.0/24
Dell(conf-redirect-list)#redirect tunnel 1 track 1 tcp any any
Dell(conf-redirect-list)#redirect tunnel 1 track 1 udp 155.55.0.0/16 host
144.144.144.144
Dell(conf-redirect-list)#redirect tunnel 2 track 2 tcp 155.55.2.0/24
222.22.2.0/24
Dell(conf-redirect-list)#redirect tunnel 2 track 2 tcp any any
Dell(conf-redirect-list)#end
Dell#
```

Apply the Redirect Rule to an Interface:

```
Dell#configure terminal
Dell(conf)#interface TenGigabitEthernet 2/28
Dell(conf-if-te-2/28)#ip redirect-group explicit_tunnel
Dell(conf-if-te-2/28)#exit
Dell(conf)#end
```

Verify the Applied Redirect Rules:

```
Dell#show ip redirect-list explicit_tunnel
```

```
IP redirect-list explicit_tunnel:
 Defined as:
 seq 5 redirect tunnel 1 track 1 tcp 155.55.2.0/24 222.22.2.0/24, Track 1
[up], Next-hop reachable (via Te 1/32)
 seq 10 redirect tunnel 1 track 1 tcp any any, Track 1 [up], Next-hop
reachable (via Te 1/32)
 seq 15 redirect tunnel 1 track 1 udp 155.55.0.0/16 host 144.144.144.144,
Track 1 [up], Next-hop reachable (via Te 1/32)
 seq 20 redirect tunnel 2 track 2 tcp 155.55.2.0/24 222.22.2.0/24, Track 2
[up], Next-hop reachable (via Te 1/33)
 seq 25 redirect tunnel 2 track 2 tcp any any, Track 2 [up], Next-hop
reachable (via Te 1/33)
 Applied interfaces:
```

Te 2/28  
Dell#

## PIM Sparse-Mode (PIM-SM)

Protocol-independent multicast sparse-mode (PIM-SM) is supported on the MXL switch platform.

PIM-SM is a multicast protocol that forwards multicast traffic to a subnet only after a request using a PIM Join message; this behavior is the opposite of PIM-Dense mode, which forwards multicast traffic to all subnets until a request to stop.

### Implementation Information

Be aware of the following PIM-SM implementation information.

- The Dell Networking implementation of PIM-SM is based on IETF *Internet Draft draft-ietf-pim-sm-v2-new-05*.
- MXL supports a maximum of 31 PIM interfaces and 2K multicast entries including (\*,G), and (S,G) entries. There is no limit on the number of PIM neighbors MXL can have.
- The SPT-Threshold is zero, which means that the last-hop designated router (DR) joins the shortest path tree (SPT) to the source after receiving the first multicast packet.
- The Dell Networking operating system (OS) reduces the number of control messages sent between multicast routers by bundling Join and Prune requests in the same message.
- The Dell Networking OS supports PIM-SM on physical, virtual local area network (VLAN), and port-channel interfaces.
- The Dell Networking OS supports 2000 IPv6 multicast forwarding entries, with up to 128 PIM-source-specific multicast (SSM) neighbors/interfaces.
- IPv6 Multicast is not supported on synchronous optical network technologies (SONET) interfaces.

### Protocol Overview

PIM-SM initially uses unidirectional shared trees to forward multicast traffic; that is, all multicast traffic must flow only from the rendezvous point (RP) to the receivers.

After a receiver receives traffic from the RP, PM-SM switches to SPT to forward multicast traffic. Every multicast group has an RP and a unidirectional shared tree (group-specific shared tree).

### Requesting Multicast Traffic

A host requesting multicast traffic for a particular group sends an Internet group management protocol (IGMP) Join message to its gateway router.

The gateway router is then responsible for joining the shared tree to the RP (RPT) so that the host can receive the requested traffic.

1. After receiving an IGMP Join message, the receiver gateway router (last-hop DR) creates a (\*,G) entry in its multicast routing table for the requested group. The interface on which the join message was received becomes the outgoing interface associated with the (\*,G) entry.
2. The last-hop DR sends a PIM Join message to the RP. All routers along the way, including the RP, create an (\*,G) entry in their multicast routing table, and the interface on which the message was received becomes the outgoing interface associated with the (\*,G) entry. This process constructs an RPT branch to the RP.
3. If a host on the same subnet as another multicast receiver sends an IGMP report for the same multicast group, the gateway takes no action. If a router between the host and the RP receives a PIM Join message for which it already has a (\*,G) entry, the interface on which the message was received is added to the outgoing interface list associated with the (\*,G) entry, and the message is not (and does not need to be) forwarded towards the RP.

## Refuse Multicast Traffic

A host requesting to leave a multicast group sends an IGMP Leave message to the last-hop DR. If the host is the only remaining receiver for that group on the subnet, the last-hop DR is responsible for sending a PIM Prune message up the RPT to prune its branch to the RP.

1. After receiving an IGMP Leave message, the gateway removes the interface on which it is received from the outgoing interface list of the (\*,G) entry. If the (\*,G) entry has no remaining outgoing interfaces, multicast traffic for that group is no longer forwarded to that subnet.
2. If the (\*,G) entry has no remaining outgoing interfaces, the last-hop DR sends a PIM Prune message to towards the RP. All routers along the way remove the interface on which the message was received from the outgoing interface list of the (\*,G) entry. If on any router there is at least one outgoing interface listed for that (\*,G) entry, the Prune message is not forwarded.

## Send Multicast Traffic

With PIM-SM, all multicast traffic must initially originate from the RP. A source must unicast traffic to the RP so that the RP can learn about the source and create an SPT to it. Then the last-hop DR may create an SPT directly to the source.

1. The source gateway router (first-hop DR) receives the multicast packets and creates an (S,G) entry in its multicast routing table. The first-hop DR encapsulates the initial multicast packets in PIM Register packets and unicasts them to the RP.
2. The RP decapsulates the PIM Register packets and forwards them if there are any receivers for that group. The RP sends a PIM Join message towards the source. All routers between the RP and the source, including the RP, create an (S,G) entry and list the interface on which the message was received as an outgoing interface, thus recreating a SPT to the source.
3. After the RP starts receiving multicast traffic via the (S,G), it unicasts a Register-Stop message to the first-hop DR so that multicast packets are no longer encapsulated in PIM Register packets and unicast. After receiving the first multicast packet from a particular source, the last-hop DR sends a PIM Join message to the source to create an SPT to it.
4. There are two paths, then, between the receiver and the source, a direct SPT and an RPT. One router receives a multicast packet on two interfaces from the same source in this case; this router prunes the shared tree by sending a PIM Prune message to the RP that tells all routers between the source and the RP to remove the outgoing interface from the (\*,G) entry, and tells the RP to prune its SPT to the source with a Prune message.

**Dell Networking Behavior:** When the router creates an SPT to the source, there are then two paths between the receiver and the source, the SPT and the RPT. Until the router can prune itself from the RPT, the receiver receives duplicate multicast packets which may cause disruption. Therefore, the router must prune itself from the RPT as soon as possible. The Dell Networking OS optimizes the shared to shortest-

path tree switchover latency by copying and forwarding the first (S,G) packet received on the SPT to the PIM task immediately upon arrival. The arrival of the (S,G) packet confirms for PIM that the SPT is created, and that it can prune itself from the shared tree.

### Important Point to Remember

If you use a Loopback interface with a /32 mask as the RP, you must enable PIM Sparse-mode on the interface.

## Configuring PIM-SM

Configuring PIM-SM is a three-step process.

1. Enable multicast routing (refer to the following step).
2. Select a rendezvous point.
3. Enable PIM-SM on an interface.

Enable multicast routing.

CONFIGURATION mode

```
ip multicast-routing
```

### Related Configuration Tasks

The following are related PIM-SM configuration tasks.

- [Configuring S,G Expiry Timers](#)
- [Configuring a Static Rendezvous Point](#)
- [Configuring a Designated Router](#)
- [Creating Multicast Boundaries and Domains](#)

## Enable PIM-SM

You must enable PIM-SM on each participating interface.

1. Enable multicast routing on the system.

CONFIGURATION mode

```
ip multicast-routing
```

2. Enable PIM-Sparse mode.

INTERFACE mode

```
ip pim sparse-mode
```


## Example of Viewing PIM-SM Interfaces

### Example of Viewing PIM Neighbors

### Example of Viewing the PIM Multicast Routing Table

To display which interfaces are enabled with PIM-SM, use the `show ip pim interface` command from EXEC Privilege mode.

```
Dell#show ip pim interface
Address Interface VIFindex Ver/ Nbr Query DR DR
 Mode Count Intvl Prio
189.87.5.6 Gi 4/11 0x2 v2/S 1 30 1 127.87.5.6
189.87.3.2 Gi 4/12 0x3 v2/S 1 30 1 127.87.3.5
189.87.31.6 Gi 7/11 0x0 v2/S 0 30 1 127.87.31.6
189.87.50.6 Gi 7/13 0x4 v2/S 1 30 1 127.87.50.6
Dell#
```

 **NOTE:** You can influence the selection of the Rendezvous Point by enabling PIM-Sparse mode on a Loopback interface and assigning a low IP address.

To display PIM neighbors for each interface, use the `show ip pim neighbor` command EXEC Privilege mode.

```
Dell#show ip pim neighbor
Neighbor Interface Uptime/Expires Ver DR
Address
127.87.5.5 Gi 4/11 01:44:59/00:01:16 v2 1 / S
127.87.3.5 Gi 4/12 01:45:00/00:01:16 v2 1 / DR
127.87.50.5 Gi 7/13 00:03:08/00:01:37 v2 1 / S
Dell#
```

To display the PIM routing table, use the `show ip pim rib` command from EXEC privilege mode.

```
Dell#show ip pim rib

PIM Multicast Routing Table
Flags: D - Dense, S - Sparse, C - Connected, L - Local, P - Pruned,
 R - RP-bit set, F - Register flag, T - SPT-bit set, J - Join SPT,
Timers: Uptime/Expires
Interface state: Interface, next-Hop, State/Mode

(*, 192.1.2.1), uptime 00:29:36, expires 00:03:26, RP 10.87.2.6, flags: SCJ
 Incoming interface: GigabitEthernet 4/12, RPF neighbor 10.87.3.5
 Outgoing interface list:
 GigabitEthernet 4/11
 GigabitEthernet 7/13

(10.87.31.5, 192.1.2.1), uptime 00:01:24, expires 00:02:26, flags: FT
 Incoming interface: GigabitEthernet 7/11, RPF neighbor 0.0.0.0
 Outgoing interface list:
 GigabitEthernet 4/11
 GigabitEthernet 4/12
 GigabitEthernet 7/13
--More--
```

## Configuring S,G Expiry Timers

By default, S, G entries expire in 210 seconds. You can configure a global expiry time (for all [S,G] entries) or configure an expiry time for a particular entry.

If you configure both, the ACL supersedes the global configuration for the specified entries.

When you create, delete, or update an expiry time, the changes are applied when the keep alive timer refreshes.

To configure a global expiry time or to configure the expiry time for a particular (S,G) entry, use the following commands.

1. Enable global expiry timer for S, G entries.

CONFIGURATION mode

```
ip pim sparse-mode sg-expiry-timer seconds
```

The range is from 211 to 86,400 seconds.

The default is **210**.

2. Create an extended ACL.

CONFIGURATION mode

```
ip access-list extended access-list-name
```

3. Specify the source and group to which the timer is applied using extended ACLs with permit rules only.

CONFIG-EXT-NAACL mode

```
[seq sequence-number] permit ip source-address/mask | any | host source-address {destination-address/mask | any | host destination-address}
```

4. Set the expiry time for a specific (S,G) entry (as shown in the following example).

CONFIGURATION mode

```
ip pim sparse-mode sg-expiry-timer seconds sg-list access-list-name
```

The range is from 211 to 86,400 seconds.

The default is **210**.

### Example Configuring an (S,G) Expiry Time

 **NOTE:** The expiry time configuration is nullified and the default global expiry time is used if:

- an ACL is specified in the `ip pim sparse-mode sg-expiry-timer` command, but the ACL has not been created or is a standard ACL.
- if the expiry time is specified for an (S,G) entry in a deny rule.

```
Dell(conf)#ip access-list extended SGtimer
Dell(config-ext-nacl)#permit ip 10.1.2.3/24 225.1.1.0/24
Dell(config-ext-nacl)#permit ip any 232.1.1.0/24
Dell(config-ext-nacl)#permit ip 100.1.1.0/16 any
Dell(config-ext-nacl)#show conf
```

```

!
ip access-list extended SGtimer
 seq 5 permit ip 10.1.2.0/24 225.1.1.0/24
 seq 10 permit ip any 232.1.1.0/24
 seq 15 permit ip 100.1.0.0/16 any
Dell(config-ext-nacl)#exit
Dell(conf)#ip pim sparse-mode sg-expiry-timer 1800 sg-list SGtimer

```

To display the expiry time configuration, use the `show running-configuration [acl | pim]` command from EXEC Privilege mode.

## Configuring a Static Rendezvous Point

The rendezvous point (RP) is a PIM-enabled interface on a router that acts as the root a group-specific tree; every group must have an RP.

- Identify an RP by the IP address of a PIM-enabled or Loopback interface.
 

```
ip pim rp-address
```

### Example of Viewing an RP on a Loopback Interface

```

Dell#sh run int loop0
!
interface Loopback 0
 ip address 1.1.1.1/32
 ip pim sparse-mode
 no shutdown
Dell#sh run pim
!
ip pim rp-address 1.1.1.1 group-address 224.0.0.0/4

```

## Overriding Bootstrap Router Updates

PIM-SM routers must know the address of the RP for each group for which they have (\*,G) entry. This address is obtained automatically through the bootstrap router (BSR) mechanism or a static RP configuration.

Use the following command if you have configured a static RP for a group. If you do not use the `override` option with the following command, the RPs advertised in the BSR updates take precedence over any statically configured RPs.

- Use the `override` option to override bootstrap router updates with your static RP configuration.
 

```
ip pim rp-address
```

### Example of Viewing the Rendezvous Point (Multicast Group)

### Example of Viewing the Rendezvous Point (Multicast Group Range)

To display the assigned RP for a group, use the `show ip pim rp` command from EXEC privilege mode.

```

Dell#show ip pim rp
Group RP
225.0.1.40 165.87.50.5
226.1.1.1 165.87.50.5

```



To display the assigned RP for a group range (group-to-RP mapping), use the `show ip pim rp mapping` command in EXEC privilege mode.

```
Dell#show ip pim rp mapping
PIM Group-to-RP Mappings
Group(s): 224.0.0.0/4, Static
 RP: 165.87.50.5, v2
```

## Configuring a Designated Router

Multiple PIM-SM routers might be connected to a single local area network (LAN) segment. One of these routers is elected to act on behalf of directly connected hosts. This router is the designated router (DR). The DR is elected using hello messages. Each PIM router learns about its neighbors by periodically sending a hello message out of each PIM-enabled interface. Hello messages contain the IP address of the interface out of which it is sent and a DR priority value. The router with the greatest priority value is the DR. If the priority value is the same for two routers, then the router with the greatest IP address is the DR. By default, the DR priority value is 192, so the IP address determines the DR.

- Assign a DR priority value.

INTERFACE mode

```
ip pim dr-priority priority-value
```

- Change the interval at which a router sends hello messages.

INTERFACE mode

```
ip pim query-interval seconds
```

- Display the current value of these parameter.

EXEC Privilege mode

```
show ip pim interface
```

## Creating Multicast Boundaries and Domains

A PIM domain is a contiguous set of routers that all implement PIM and are configured to operate within a common boundary defined by PIM multicast border routers (PMBRs). PMBRs connect each PIM domain to the rest of the Internet.

Create multicast boundaries and domains by filtering inbound and outbound bootstrap router (BSR) messages per interface. The following command is applied to the subsequent inbound and outbound updates. Timeout removes existing BSR advertisements.

- Create multicast boundaries and domains by filtering inbound and outbound BSR messages per interface.

```
ip pim bsr-border
```

- Remove candidate RP advertisements.

```
clear ip pim rp-mapping
```

## Enabling PIM-SM Graceful Restart

To enable PIM-SM graceful restart, use the following commands.

- Enable PIM-SM graceful restart (non-stop forwarding capability).  
CONFIGURATION mode

```
ip pim graceful-restart nsf
```

- (option) `restart-time`: the time the Dell Networking system requires to restart. The default value is **180 seconds**.
- (option) `stale-entry-time`: the maximum amount of time that the Dell Networking system preserves entries from a restarting neighbor. The default value is **60 seconds**.
- (option) `helper-only`: this mode takes precedence over any graceful restart configuration.

## PIM Source-Specific Mode (PIM-SSM)

PIM source-specific mode (PIM-SSM) is supported on the MXL switch platform.

PIM-SSM is a multicast protocol that forwards multicast traffic from a single source to a subnet. In the other versions of protocol independent multicast (PIM), a receiver subscribes to a group only. The receiver receives traffic not just from the source in which it is interested but from all sources sending to that group. PIM-SSM requires that receivers specify the sources in which they are interested using IGMPv3 include messages to avoid receiving unwanted traffic.

PIM-SSM is more efficient than PIM-SM because it immediately creates shortest path trees (SPT) to the source rather than first using shared trees. PIM-SM requires a shared tree rooted at the RP because IGMPv2 receivers do not know about the source sending multicast data. Multicast traffic passes from the source to the receiver through the RP, until the receiver learns the source address, at which point it switches to the SPT. PIM-SSM uses IGMPv3. Because receivers subscribe to a source and group, the RP and shared tree is unnecessary; only SPTs are used. On Dell Networking systems, it is possible to use PIM-SM with IGMPv3 to achieve the same result, but PIM-SSM eliminates the unnecessary protocol overhead.

PIM-SSM also solves the multicast address allocation problem. Applications must use unique multicast addresses because if multiple applications use the same address, receivers receive unwanted traffic. However, global multicast address space is limited. Currently GLOP/EGLOP is used to statically assign Internet-routable multicast addresses, but each autonomous system number yields only 255 multicast addresses. For short-term applications, an address could be leased, but no global dynamic multicast address allocation scheme has been accepted yet. PIM-SSM eliminates the need for unique multicast addresses because routing decisions for (S1, G1) are independent from (S2, G1). As a result, subnets do not receive unwanted traffic when multiple applications use the same address.

## Configure PIM-SMM

Configuring PIM-SSM is a two-step process.

1. Configure PIM-SMM.
2. Enable PIM-SSM for a range of addresses.

### Related Configuration Tasks

- [Use PIM-SSM with IGMP Version 2 Hosts](#)

## Implementation Information

- The Dell Networking implementation of PIM-SSM is based on RFC 3569.
- The Dell Networking operating system (OS) reduces the number of control messages sent between multicast routers by bundling Join and Prune requests in the same message.

### Important Points to Remember

- The default SSM range is 232/8 always. Applying an SSM range does not overwrite the default range. Both the default range and SSM range are effective even when the default range is not added to the SSM ACL.
- Extended ACLs cannot be used for configuring SSM range. Be sure to create the ACL first and then apply it to the SSM range.
- The default range is always supported, so range can never be smaller than the default.

## Enabling PIM-SSM

To enable PIM-SSM, follow these steps.

1. Create an ACL that uses permit rules to specify what range of addresses should use SSM.

```
CONFIGURATION mode
```

```
ip access-list standard name
```

2. Enter the `ip pim ssm-range` command and specify the ACL you created.

```
CONFIGURATION mode
```

```
ip pim ssm-range acl-name
```

### Enabling PIM-SSM

To display address ranges in the PIM-SSM range, use the `show ip pim ssm-range` command from EXEC Privilege mode.

```
R1(conf)#do show run pim
!
ip pim rp-address 10.11.12.2 group-address 224.0.0.0/4
ip pim ssm-range ssm
R1(conf)#do show run acl
!
ip access-list standard ssm
 seq 5 permit host 239.0.0.2
R1(conf)#do show ip pim ssm-range
Group Address / MaskLen
239.0.0.2 / 32
```

## Use PIM-SSM with IGMP Version 2 Hosts

PIM-SSM requires receivers that support IGMP version 3. You can employ PIM-SSM even when receivers support only IGMP version 1 or version 2 by translating (\*,G) entries to (S,G) entries.

Translate (\*,G) entries to (S,G) entries using the `ip igmp ssm-map acl` command source from CONFIGURATION mode. In a standard access list, specify the groups or the group ranges that you want to map to a source. Then, specify the multicast source.

- When an SSM map is in place and the Dell Networking OS cannot find any matching access lists for a group, it continues to create (\*,G) entries because there is an implicit deny for unspecified groups in the ACL.
- When you remove the mapping configuration, the Dell Networking OS removes the corresponding (S,G) states that it created and re-establishes the original (\*,G) states.
- You may enter multiple `ssm-map` commands for different access lists. You may also enter multiple `ssm-map` commands for the same access list, as long as they use different source addresses.
- When an extended ACL is associated with this command, the system displays an error message. If you apply an extended ACL before you create it, the system accepts the configuration, but when the ACL is later defined, the system ignores the ACL and the stated mapping has no effect.

To display the source to which a group is mapped, use the `show ip igmp ssm-map [group]` command. If you use the `group` option, the command displays the group-to-source mapping even if the group is not currently in the IGMP group table. If you do not specify the `group` option, the display is a list of groups currently in the IGMP group table that has a group-to-source mapping.

To display the list of sources mapped to a group currently in the IGMP group table, use the `show ip igmp groups group detail` command.

## Configuring PIM-SSM with IGMPv2

```
R1(conf)#do show run pim
!
ip pim rp-address 10.11.12.2 group-address 224.0.0.0/4
ip pim ssm-range ssm
R1(conf)#do show run acl
!
ip access-list standard map
seq 5 permit host 239.0.0.2
!
ip access-list standard ssm
seq 5 permit host 239.0.0.2
R1(conf)#ip igmp ssm-map map 10.11.5.2
R1(conf)#do show ip igmp groups
Total Number of Groups: 2
IGMP Connected Group Membership
Group Address Interface Mode Uptime Expires Last Reporter
239.0.0.2 Vlan 300 IGMPv2-Compat 00:00:07 Never 10.11.3.2
 Member Ports: Gi 1/1
239.0.0.1 Vlan 400 INCLUDE 00:00:10 Never 10.11.4.2
R1(conf)#do show ip igmp ssm-map
IGMP Connected Group Membership
Group Address Interface Mode Uptime Expires Last Reporter
239.0.0.2 Vlan 300 IGMPv2-Compat 00:00:36 Never 10.11.3.2
 Member Ports: Gi 1/1
R1(conf)#do show ip igmp ssm-map 239.0.0.2
```

```
SSM Map Information
Group : 239.0.0.2
Source(s) : 10.11.5.2
R1(conf)#do show ip igmp groups detail
```

```
Interface Vlan 300
Group 239.0.0.2
Uptime 00:00:01
Expires Never
Router mode IGMPv2-Compat
Last reporter 10.11.3.2
Last reporter mode IGMPv2
Last report received Join
Group source list
Source address Uptime Expires
10.11.5.2 00:00:01 Never
```

```
Interface Vlan 400
Group 239.0.0.1
Uptime 00:00:05
Expires Never
Router mode INCLUDE
Last reporter 10.11.4.2
Last reporter mode INCLUDE
Last report received ALLOW
Group source list
Source address Uptime Expires
10.11.5.2 00:00:05 00:02:04
Member Ports: Gi 1/2
```

## Port Monitoring

Port monitoring is supported on the MXL switch platform.

Mirroring is used for monitoring Ingress or Egress or both Ingress and Egress traffic on a specific port(s). This mirrored traffic can be sent to a port where a network sniffer can connect and monitor the traffic.

Dell Networking OS supports the following mirroring techniques:

- Port-Mirroring — Port Monitoring is a method of monitoring network traffic that forwards a copy of each incoming and outgoing packet from one port of a network router to another port where the packet can be studied.
- Remote Port Monitoring (RPM) — Remote Port Monitoring allows the user to monitor traffic running across a remote device in the same network. Here the mirror traffic is carried over the L2 network, so that probe devices in the network can analyze it. It is an extension to the normal Port Monitoring feature. This feature is generally referred as RPM, where mirror traffic is carried over L2 network.
- Encapsulated Remote-Port Monitoring (ERPM) — ERPM is a feature to encapsulate mirrored packet using GRE with IP delivery so that it can be sent across a routed network.

## Important Points to Remember

- Port Monitoring is supported on both physical and logical interfaces like virtual area network (VLAN) and port-channel.
- The monitored (the source, [MD]) and monitoring ports (the destination, [MG]) must be on the same switch.
- In general, a monitoring port should have no ip address and no shutdown as the only configuration; Dell Networking OS permits a limited set of commands for monitoring ports. You can display these commands using the ? command. A monitoring port also may not be a member of a VLAN.
- There may only be one destination port (MG) in a monitoring session.
- Source port (MD) can be monitored by more than one destination port (MG).
- Destination port (MG) can be a physical interface or port-channel interface.
- A Port monitoring session can have multiple source statements.
- Range command is supported in the source statement, where we can specify a range of interfaces of (Physical, Port Channel or VLAN) types.
- One Destination Port (MG) can be used in multiple sessions.
- There can be a maximum of 128 source ports in a Port Monitoring session.
- Flow based monitoring is supported for all type of source interfaces.
- Source port (MD) can be a VLAN, where the VLAN traffic received on that port pipe where its members are present is monitored
- Single MD can be monitored on max. of 4 MG ports. If you try to assign a monitored port to more than one monitoring port, the following message displays:

```
Dell (conf-mon-sess-2)#do show monitor session
 SessID Source Destination Dir Mode Source IP Dest IP
```

```

1 Te 0/0 Te 0/1 both Port N/A N/A
2 Te 0/0 Te 0/2 both Port N/A N/A
Dell (conf-mon-sess-2)#do show running-config monitor session
!
monitor session 1
 source TenGigabitEthernet 0/0 destination TenGigabitEthernet 0/1 direction both
!
monitor session 2
 source TenGigabitEthernet 0/0 destination TenGigabitEthernet 0/2 direction both
Dell (conf-mon-sess-2)#
!

```

## Configuring Port Monitoring

To configure port monitoring, use the following commands.

1. Verify that the intended monitoring port has no configuration other than no shutdown, as shown in the following example.

EXEC Privilege mode

```
show interface
```

2. Create a monitoring session using the command `monitor session` from CONFIGURATION mode, as shown in the following example.

CONFIGURATION mode

```
monitor session
```

```
monitor session type rpm/erpm
type is an optional keyword, required only for rpm and erpm
```

3. Specify the source and destination port and direction of traffic, as shown in the following example.

MONITOR SESSION mode

```
source
```

### Example of Viewing Port Monitoring Configuration

To display information on currently configured port-monitoring sessions, use the `show monitor session` command from EXEC Privilege mode.

```

Dell(conf)#monitor session 0
Dell(conf-mon-sess-0)#$source ten 0/0 dest ten 0/1 dir rx
Dell(conf-mon-sess-0)#show c
!
monitor session 0
 source TenGigabitEthernet 0/0 destination TenGigabitEthernet 0/1 direction rx
Dell(conf-mon-sess-0)#
Dell(conf-mon-sess-0)#do show monitor session
 SessID Source Destination Dir Mode Source IP Dest IP

 0 Te 0/0 Te 0/1 rx Port N/A N/A

Dell(conf)#monitor session 0
Dell(conf-mon-sess-0)#source po 10 dest ten 0/1 dir rx
Dell(conf-mon-sess-0)#do show monitor session
 SessID Source Destination Dir Mode Source IP Dest IP

 0 Te 0/0 Te 0/1 rx Port N/A N/A

```



```

0 Po 10 Te 0/1 rx Port N/A N/A

Dell(conf)#monitor session 1
Dell(conf-mon-sess-1)#source vl 40 dest ten 0/2 dir rx
Dell(conf-mon-sess-1)#flow-based enable
Dell(conf-mon-sess-1)#exit
Dell(conf)#do show monitor session

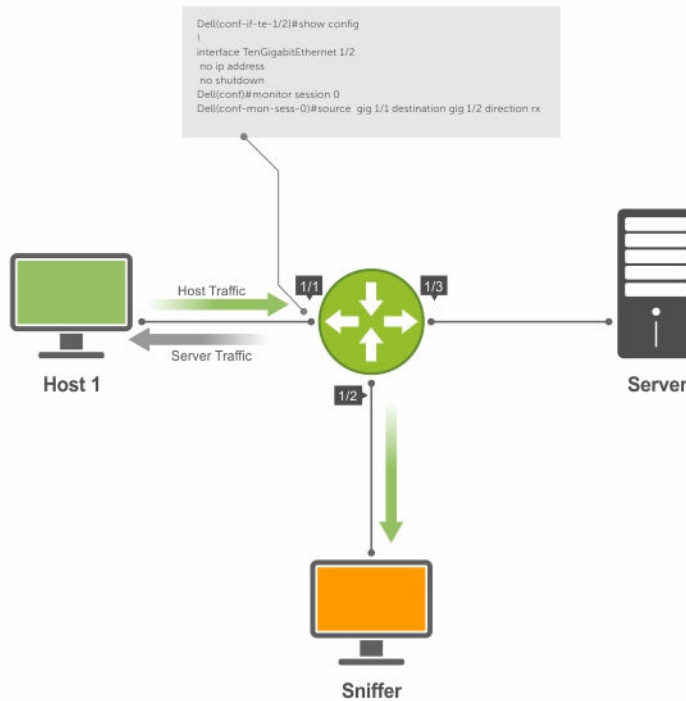
```

| SessID | Source | Destination | Dir | Mode | Source IP | Dest IP |
|--------|--------|-------------|-----|------|-----------|---------|
| 0      | Te 0/0 | Te 0/1      | rx  | Port | N/A       | N/A     |
| 0      | Po 10  | Te 0/1      | rx  | Port | N/A       | N/A     |
| 1      | Vl 40  | Te 0/2      | rx  | Flow | N/A       | N/A     |

Note: Source as VLAN is achieved via Flow based mirroring. Please refer section [Enabling Flow-Based monitoring](#).

In the following example, the host and server are exchanging traffic which passes through the uplink interface 1/1. Port 1/1 is the monitored port and port 1/42 is the destination port, which is configured to only monitor traffic received on tengigabitethernet 1/1 (host-originated traffic).

**Figure 98. Port Monitoring Example**



## Enabling Flow-Based Monitoring

Flow-based monitoring is supported only on the S-Series platform.

Flow-based monitoring conserves bandwidth by monitoring only specified traffic instead of all traffic on the interface. This feature is particularly useful when looking for malicious traffic. It is available for Layer 2 and Layer 3 ingress and egress traffic. You can specify traffic using standard or extended access-lists.

1. Enable flow-based monitoring for a monitoring session.

MONITOR SESSION mode

```
flow-based enable
```

2. Define in access-list rules that include the keyword `monitor`. For port monitoring, Dell Networking OS only considers traffic matching rules with the keyword `monitor`.

CONFIGURATION mode

```
ip access-list
```

Refer to [Access Control Lists \(ACLs\)](#).

3. Apply the ACL to the monitored port.

INTERFACE mode

```
ip access-group access-list
```

### Example of the `flow-based enable` Command

To view an access-list that you applied to an interface, use the `show ip accounting access-list` command from EXEC Privilege mode.

```
Dell(conf)#monitor session 0
Dell(conf-mon-sess-0)#flow-based enable
Dell(conf)#ip access-list ext testflow
Dell(config-ext-nacl)#seq 5 permit icmp any any count bytes monitor
Dell(config-ext-nacl)#seq 10 permit ip 102.1.1.0/24 any count bytes monitor
Dell(config-ext-nacl)#seq 15 deny udp any any count bytes
Dell(config-ext-nacl)#seq 20 deny tcp any any count bytes
Dell(config-ext-nacl)#exit
Dell(conf)#interface gig 1/1
Dell(conf-if-gi-1/1)#ip access-group testflow in
Dell(conf-if-gi-1/1)#show config
!
interface GigabitEthernet 1/1
 ip address 10.11.1.254/24
 ip access-group testflow in
 shutdown
Dell(conf-if-gi-1/1)#exit
Dell(conf)#do show ip accounting access-list testflow
!
Extended Ingress IP access list testflow on GigabitEthernet 1/1
Total cam count 4
 seq 5 permit icmp any any monitor count bytes (0 packets 0 bytes)
 seq 10 permit ip 102.1.1.0/24 any monitor count bytes (0 packets 0 bytes)
 seq 15 deny udp any any count bytes (0 packets 0 bytes)
 seq 20 deny tcp any any count bytes (0 packets 0 bytes)
Dell(conf)#do show monitor session 0
SessionID Source Destination Direction Mode Type
```

-----  
0            Gi 1/1 Gi 1/2            rx            interface Flow-based  
-----

## Remote Port Mirroring

Remote Port Mirroring is supported on the MXL Switch platform.

While local port monitoring allows you to monitor traffic from one or more source ports by directing it to a destination port on the same switch/router, remote port mirroring allows you to monitor Layer 2 and Layer 3 ingress or egress or both ingressing or egressing traffic on multiple source ports on different switches and forward the mirrored traffic to multiple destination ports on different switches. Remote port mirroring helps network administrators monitor and analyze traffic to troubleshoot network problems in a time-saving and efficient way.

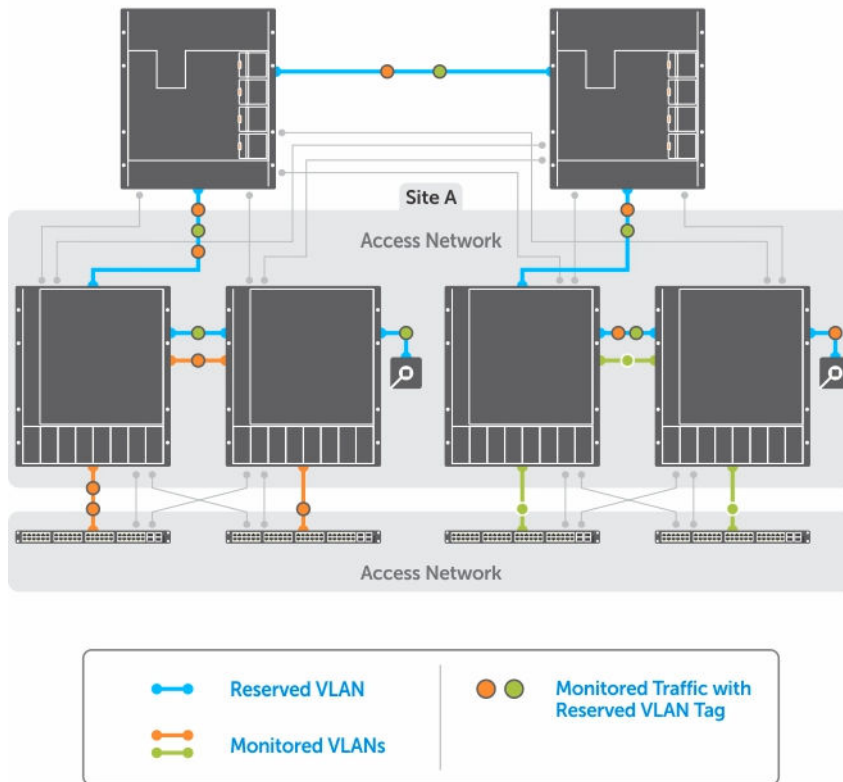
In a remote-port mirroring session, monitored traffic is tagged with a VLAN ID and switched on a user-defined, non-routable L2 VLAN. The VLAN is reserved in the network to carry only mirrored traffic, which is forwarded on all egress ports of the VLAN. Each intermediate switch that participates in the transport of mirrored traffic must be configured with the reserved L2 VLAN. Remote port monitoring supports mirroring sessions in which multiple source and destination ports are distributed across multiple switches

### Remote Port Mirroring Example

Remote port mirroring uses the analyzers shown in the aggregation network in Site A.

The VLAN traffic on monitored links from the access network is tagged and assigned to a dedicated L2 VLAN. Monitored links are configured in two source sessions shown with orange and green circles. Each source session uses a separate reserved VLAN to transmit mirrored packets (mirrored source-session traffic is shown with an orange or green circle with a blue border).

The reserved VLANs transport the mirrored traffic in sessions (blue pipes) to the destination analyzers in the local network. Two destination sessions are shown: one for the reserved VLAN that transports orange-circle traffic; one for the reserved VLAN that transports green-circle traffic.



## Configuring Remote Port Mirroring

Remote port mirroring requires a source session (monitored ports on different source switches), a reserved tagged VLAN for transporting mirrored traffic (configured on source, intermediate, and destination switches), and a destination session (destination ports connected to analyzers on destination switches).

### Configuration Notes

When you configure remote port mirroring, the following conditions apply:

- You can configure any switch in the network with source ports and destination ports, and allow it to function in an intermediate transport session for a reserved VLAN at the same time for multiple remote-port mirroring sessions. You can enable and disable individual mirroring sessions.
- BPDU monitoring is not required to use remote port mirroring.
- A remote port mirroring session mirrors monitored traffic by prefixing the reserved VLAN tag to monitored packets so that they are copied to the reserve VLAN.
- Mirrored traffic is transported across the network using 802.1Q-in-802.1Q tunneling. The source address, destination address and original VLAN ID of the mirrored packet are preserved with the tagged VLAN header. Untagged source packets are tagged with the reserve VLAN ID.
- The RPM VLAN can't be a Private VLAN.
- The RPM VLAN can be used as GVRP VLAN.

- The L3 interface configuration should be blocked for RPM VLAN.
- The member port of the reserved VLAN should have MTU and IPMTU value as MAX+4 (to hold the VLAN tag parameter).
- To associate with source session, the reserved VLAN can have at max of only 4 member ports.
- To associate with destination session, the reserved VLAN can have multiple member ports.
- Reserved Vlan cannot have untagged ports.

In the reserved **L2 VLAN** used for remote port mirroring:

- MAC address learning in the reserved VLAN is automatically disabled.
- The reserved VLAN for remote port mirroring can be automatically configured in intermediate switches by using GVRP.
- There is no restriction on the VLAN IDs used for the reserved remote-mirroring VLAN. Valid VLAN IDs are from 2 to 4094. The default VLAN ID is not supported.
- In mirrored traffic, packets that have the same destination MAC address as an intermediate or destination switch in the path used by the reserved VLAN to transport the mirrored traffic are dropped by the switch that receives the traffic if the switch has a L3 VLAN configured.

In a **source session** used for remote port mirroring:

- Maximum number of source sessions supported on a switch: 4
- Maximum number of source ports supported in a source session: 128
- You can configure physical ports and port-channels as sources in remote port mirroring and use them in the same source session. You can use both Layer 2 (configured with the switchport command) and Layer 3 ports as source ports. You can optionally configure one or more source VLANs to specify the VLAN traffic to be mirrored on source ports.
- You can use the default VLAN and native VLANs as a source VLAN.
- You cannot configure the dedicated VLAN used to transport mirrored traffic as a source VLAN.
- Egressing remote-vlan packets are rate limited to a default value of 100 Mbps.

In a **destination session** used for remote port mirroring:

- Maximum number of destination sessions supported on a switch: 64
- Maximum number ports supported in a destination session: 64.
- You can configure any port as a destination port.
- You can configure additional destination ports in an active session.
- You can tunnel the mirrored traffic from multiple remote-port source sessions to the same destination port.
- By default, destination port sends the mirror traffic to the probe port by stripping off the rpm header. We can also configure the destination port to send the mirror traffic with the rpm header intact in the original mirror traffic.
- By default, ingress traffic on a destination port is dropped.

## Restrictions

When you configure remote port mirroring, the following **restrictions** apply:

- You can configure the same source port to be used in multiple source sessions.
- You cannot configure a source port channel or source VLAN in a source session if the port channel or VLAN has a member port that is configured as a destination port in a remote-port mirroring session.

- A destination port for remote port mirroring cannot be used as a source port, including the session in which the port functions as the destination port.
- A destination port cannot be used in any spanning tree instance.
- The reserved VLAN used to transport mirrored traffic must be a L2 VLAN. L3 VLANs are not supported.
- On a source switch on which you configure source ports for remote port mirroring, you can add only one port to the dedicated RPM VLAN which is used to transport mirrored traffic. You can configure multiple ports for the dedicated RPM VLAN on intermediate and destination switches.

## Displaying Remote-Port Mirroring Configurations

To display the current configuration of remote port mirroring for a specified session, enter the **show config** command in **MONITOR SESSION** configuration mode.

```
Dell(conf-mon-sess-2)#show config
!
monitor session 2 type rpm
 source fortyGigE 0/60 destination remote-vlan 300 direction rx
 source Port-channel 10 destination remote-vlan 300 direction rx
 no disable
```

To display the currently configured source and destination sessions for remote port mirroring on a switch, enter the **show monitor session** command in **EXEC** Privilege mode.

```
Dell(conf)#do show monitor session
```

| SessID | Source          | Destination     | Dir | Mode | Source IP | Dest IP |
|--------|-----------------|-----------------|-----|------|-----------|---------|
| 1      | remote-vlan 100 | Fo 0/48         | N/A | N/A  | N/A       | N/A     |
| 1      | remote-vlan 100 | Po 100          | N/A | N/A  | N/A       | N/A     |
| 2      | Fo 0/60         | remote-vlan 300 | rx  | Port | N/A       | N/A     |
| 2      | Po 10           | remote-vlan 300 | rx  | Port | N/A       | N/A     |

To display the current configuration of the reserved VLAN, enter the **show vlan** command.

```
Dell#show vlan
```

```
Codes: * - Default VLAN, G - GVRP VLANs, R - Remote Port Mirroring VLANs, P -
Primary, C - Community, I - Isolated
O - Openflow
```

```
Q: U - Untagged, T - Tagged
x - Dot1x untagged, X - Dot1x tagged
o - OpenFlow untagged, O - OpenFlow tagged
G - GVRP tagged, M - Vlan-stack
i - Internal untagged, I - Internal tagged, v - VLT untagged, V - VLT tagged
```

| NUM | Status | Description | Q Ports   |
|-----|--------|-------------|-----------|
| *   | 1      | Inactive    |           |
| R   | 100    | Active      | T Fo 0/44 |
| R   | 300    | Active      | T Fo 0/52 |

## Configuring the Sample Remote Port Mirroring

Remote port mirroring requires a source session (monitored ports on different source switches), a reserved tagged VLAN for transporting mirrored traffic (configured on source, intermediate, and destination switches), and a destination session (destination ports connected to analyzers on destination switches).

### Configuration Steps for RPM

| Step | Command                                            | Purpose                                                                                                              |
|------|----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| 1    | configure terminal                                 | Enter global configuration mode.                                                                                     |
| 2    | <b>monitor session &lt;id&gt; type rpm</b>         | The <id> needs to be unique and not already defined in the box specifying type as 'rpm' defines a RPM session.       |
| 3    | <b>source</b> <i>Interface   Range</i>             | Specify the port or list of ports that needs to be monitored                                                         |
| 4    | <b>direction</b>                                   | Specify rx, tx or both in case to monitor ingress/ egress or both ingress and egress packets on the specified port.. |
| 5    | <b>rpm source-ip &lt;id&gt; dest-ip &lt;id&gt;</b> | Specify the source ip address and the destination ip where the packet needs to be sent.                              |
| 6    | <b>&lt;no&gt; flow-based enable</b>                | Specify flow-based enable for mirroring on a flow by flow basis and also for vlan as source.                         |
| 7    | <b>no enable</b>                                   | (Optional) No disable command is mandatory in order for a rpm session to be active.                                  |

### Configuring the sample Source Remote Port Mirroring

```
Dell(conf)#interface vlan 10
Dell(conf-if-vl-10)#mode remote-port-mirroring
Dell(conf-if-vl-10)#tagged te 0/4
Dell(conf-if-vl-10)#exit
```

```
Dell(conf)#monitor session 1 type rpm
Dell(conf-mon-sess-1)#source te 0/5 destination remote-vlan 10 dir rx
Dell(conf-mon-sess-1)#no disable
Dell(conf-mon-sess-1)#exit
```

```
Dell(conf)#inte vlan 100
Dell(conf-if-vl-100)#tagged te 0/7
Dell(conf-if-vl-100)#exit
```

```
Dell(conf)#interface vlan 20
Dell(conf-if-vl-20)#mode remote-port-mirroring
Dell(conf-if-vl-20)#tagged te 0/6
Dell(conf-if-vl-20)#exit
```

```
Dell(conf)#monitor session 2 type rpm
Dell(conf-mon-sess-2)#source vlan 100 destination remote-vlan 20 dir rx
Dell(conf-mon-sess-2)#no disable
Dell(conf-mon-sess-2)#flow-based enable
Dell(conf-mon-sess-2)#exit
```

```
Dell(conf)#mac access-list standard mac_acl
Dell(config-std-macl)#permit 00:00:00:00:11:22 count monitor
Dell(config-std-macl)#exit
```

```
Dell(conf)#interface vlan 100
Dell(conf-if-vl-100)#mac access-group mac_acl1 in
Dell(conf-if-vl-100)#exit
```

```
Dell(conf)#inte te 0/30
Dell(conf-if-te-0/30)#no shutdown
Dell(conf-if-te-0/30)#switchport
Dell(conf-if-te-0/30)#exit
```

```

Dell(conf)#interface vlan 30
Dell(conf-if-vl-30)#mode remote-port-mirroring
Dell(conf-if-vl-30)#tagged te 0/30
Dell(conf-if-vl-30)#exit

Dell(conf)#interface port-channel 10
Dell(conf-if-po-10)#channel-member te 0/28-29
Dell(conf-if-po-10)#no shutdown
Dell(conf-if-po-10)#exit

Dell(conf)#monitor session 3 type rpm
Dell(conf-mon-sess-3)#source port-channel 10 dest remote-vlan 30 dir both
Dell(conf-mon-sess-3)#no disable
Dell(conf-mon-sess-3)#
Dell(conf-mon-sess-3)#exit
Dell(conf)#end
Dell#

```

```

Dell#show monitor session

```

| SessID | Source | Destination    | Dir  | Mode | Source IP | Dest IP |
|--------|--------|----------------|------|------|-----------|---------|
| 1      | Te 0/5 | remote-vlan 10 | rx   | Port | N/A       | N/A     |
| 2      | Vl 100 | remote-vlan 20 | rx   | Flow | N/A       | N/A     |
| 3      | Po 10  | remote-vlan 30 | both | Port | N/A       | N/A     |

```

Dell#

```

### Configuring the sample Source Remote Port Mirroring

```

Dell(conf)#inte te 0/0
Dell(conf-if-te-0/0)#switchport
Dell(conf-if-te-0/0)#no shutdown
Dell(conf-if-te-0/0)#exit

Dell(conf)#interface te 0/1
Dell(conf-if-te-0/1)#switchport
Dell(conf-if-te-0/1)#no shutdown
Dell(conf-if-te-0/1)#exit

Dell(conf)#interface te 0/2
Dell(conf-if-te-0/2)#switchport
Dell(conf-if-te-0/2)#no shutdown
Dell(conf-if-te-0/2)#exit

Dell(conf)#inte vlan 10
Dell(conf-if-vl-10)#mode remote-port-mirroring
Dell(conf-if-vl-10)#tagged te 0/0
Dell(conf-if-vl-10)#exit

Dell(conf)#inte vlan 20
Dell(conf-if-vl-20)#mode remote-port-mirroring
Dell(conf-if-vl-20)#tagged te 0/1
Dell(conf-if-vl-20)#exit

Dell(conf)#interface vlan 30
Dell(conf-if-vl-30)#mode remote-port-mirroring
Dell(conf-if-vl-30)#tagged te 0/2
Dell(conf-if-vl-30)#exit

Dell(conf)#monitor session 1 type rpm
Dell(conf-mon-sess-1)#source remote-vlan 10 dest te 0/3
Dell(conf-mon-sess-1)#exit

```



```

Dell(conf)#monitor session 2 type rpm
Dell(conf-mon-sess-2)#source remote-vlan 20 destination te 0/4
Dell(conf-mon-sess-2)#tagged destination te 0/4
Dell(conf-mon-sess-2)#exit

Dell(conf)#monitor session 3 type rpm
Dell(conf-mon-sess-3)#source remote-vlan 30 destination te 0/5
Dell(conf-mon-sess-3)#tagged destination te 0/5
Dell(conf-mon-sess-3)#end
Dell#
Dell#show monitor session

```

| SessID | Source         | Destination | Dir | Mode | Source IP | Dest IP |
|--------|----------------|-------------|-----|------|-----------|---------|
| 1      | remote-vlan 10 | Te 0/3      | N/A | N/A  | N/A       | N/A     |
| 2      | remote-vlan 20 | Te 0/4      | N/A | N/A  | N/A       | N/A     |
| 3      | remote-vlan 30 | Te 0/5      | N/A | N/A  | N/A       | N/A     |

```

Dell#

```

## Configuring the Encapsulated Remote Port Mirroring

The ERPM session copies traffic from the source ports/lags or source VLANs and forwards the traffic using routable GRE-encapsulated packets to the destination ip address specified in the session.

### Important:

The steps to be followed for the ERPM Encapsulation :

- Dell Networking OS supports ERPM Source session only. The Encapsulated packets terminate at the destination ip or at the analyzer.
- Make sure that the destination ip is reachable via the configured ip route (static or dynamic)
- The system MTU should be configured properly to accommodate the increased size of the ERPM mirrored packet.
- The system encapsulates the complete ingress or egress data under GRE header, IP header and outer MAC header and sends it out at the next hop interface as pointed by the routing table.
- The source IP address can be any port's ip address defined in the box but it should be unique and should not be assigned to any other system in the network.
- The keyword 'flow-based enable' should have been specified in case of source as vlan or where monitoring on a per flow basis is desired.
- The keyword monitor should have been specified in the access list rules for which we need to mirror.
- The maximum number of source ports that can be defined in a session is 128.
- The system allows to configure upto 4 ERPM sessions.
- ERPM sessions do not copy locally sourced Remote-VLAN traffic from source trunk ports that carry RPM VLANs. ERPM sessions do not copy locally sourced ERPM GRE-encapsulated traffic from source ports.
- Source Vlan monitoring can be done only for ingress packets and is not supported for egress direction.


## Configuration steps for ERPM

### ERPM sample steps for configuring the Source sessions



## ERPM Behavior on a typical Dell Networking OS

The Dell Networking OS is designed to support only the Encapsulation of the data received / transmitted at the specified source port (Port A). An ERPM destination session / decapsulation of the ERPM packets at the destination Switch are not supported.



As seen in the above figure, the packets received/transmitted on Port A will be encapsulated with an IP/GRE header plus a new L2 header and sent to the destination ip address (Port D's ip address) on the sniffer. The Header that gets attached to the packet is 38 bytes long.

If the sniffer does not support IP interface, a destination switch will be needed to receive the encapsulated ERPM packet and locally mirror the whole packet to the Sniffer or a Linux Server.

### Decapsulation of ERPM packets at the Destination IP/ Analyzer

- In order to achieve the decapsulation of the original payload from the ERPM header. The below two methods are suggested :
  - a. Using Network Analyzer
    - Install any well-known Network Packet Analyzer tool which is open source and free to download.
    - Start capture of ERPM packets on the Sniffer and save it to the trace file (for example : erpmwithheader.pcap).

- The Header that gets attached to the packet is 38 bytes long. In case of a packet with L3 VLAN, it would be 42 bytes long. The original payload /original mirrored data starts from the 39<sup>th</sup> byte in a given ERPM packet. The first 38/42 bytes of the header needs to be ignored/ chopped off.
- Some tools support options to edit the capture file. We can make use of such features (for example: **editcap** ) and chop the ERPM header part and save it to a new trace file. This new file (i.e. the original mirrored packet) can be converted back into stream and fed to any egress interface.

b. Using Python script

- Either have a Linux server's ethernet port ip as the ERPM destination ip or connect the ingress interface of the server to the ERPM MirrorToPort. The analyzer should listen in the forward/ egress interface. If there is only one interface, one can choose the ingress and forward interface to be same and listen in the tx direction of the interface.
- Download/ Write a small script (for example: **erpm.py**) such that it will strip the given ERPM packet starting from the bit where GRE header ends. Basically all the bits after 0x88BE need to be removed from the packet and sent out through another interface.
- This script **erpm.zip** is available for download at the following location: [http://en.community.dell.com/techcenter/networking/m/force10\\_networking\\_scripts/20438882.aspx](http://en.community.dell.com/techcenter/networking/m/force10_networking_scripts/20438882.aspx)
- Unzip the **erpm.zip** and copy the erpm.py file to the Linux server.
- Run the python script using the following command:

```
python erpm.py -i <ingress interface> -o <egress interface>
```

erpm.py : This is the script downloaded from the script store.

**<Ingress interface>** : Specify the interface id which is connected to the mirroring port or this should be interface whose ip address has been specified as the destination ip address in the ERPM session.

**<Egress interface>** : Specify another interface on the Linux server via which the decapsulation packets can Egress. In case there is only one interface, the ingress interface itself can be specified as Egress and the analyzer can listen in the tx direction.

## Private VLANs (PVLAN)

The private VLAN (PVLAN) feature is supported on the MXL switch platform.

For syntax details about the commands described in this chapter, refer to the Private VLANs commands chapter in the *Dell Networking OS Command Line Reference Guide*.

Private VLANs extend the Dell Networking operating system (OS) security suite by providing Layer 2 isolation between ports within the same virtual local area network (VLAN). A PVLAN partitions a traditional VLAN into subdomains identified by a primary and secondary VLAN pair. Private VLANs block all traffic to isolated ports except traffic from promiscuous ports. Traffic received from an isolated port is forwarded only to promiscuous ports or trunk ports.

Example uses of PVLANS:

- A hotel can use an isolated VLAN in a PVLAN to provide Internet access for its guests, while stopping direct access between the guest ports.
- A service provider can provide Layer 2 security for customers and use the IP addresses more efficiently, by using a separate community VLAN per customer and at the same time using the same IP subnet address space for all community and isolated VLANs mapped to the same primary VLAN.
  - In more detail, community VLANs are especially useful in the service provider environment because multiple customers are likely to maintain servers that must be strictly separated in customer-specific groups. A set of servers owned by a customer could comprise a community VLAN, so that those servers could communicate with each other, and would be isolated from other customers. Another customer might have another set of servers in another community VLAN. Another customer might want an isolated VLAN, which has one or more ports that are also isolated from each other.

## Private VLAN Concepts

Review the following PVLAN concepts before you create PVLANS on your system.

The VLAN types in a PVLAN include:

- **Community VLAN** — a type of secondary VLAN in a primary VLAN:
  - Ports in a community VLAN can communicate with each other.
  - Ports in a community VLAN can communicate with all promiscuous ports in the primary VLAN.
  - A community VLAN can only contain ports configured as host.
- **Isolated VLAN** — a type of secondary VLAN in a primary VLAN:
  - Ports in an isolated VLAN cannot talk directly to each other.
  - Ports in an isolated VLAN can only communicate with promiscuous ports in the primary VLAN.
  - An isolated VLAN can only contain ports configured as host.

- **Primary VLAN** — the base VLAN of a PVLAN:
  - A switch can have one or more primary VLANs, and it can have none.
  - A primary VLAN has one or more secondary VLANs.
  - A primary VLAN and each of its secondary VLANs decrement the available number of VLAN IDs in the switch.
  - A primary VLAN has one or more promiscuous ports.
  - A primary VLAN might have one or more trunk ports, or none.
- **Secondary VLAN** — a subdomain of the primary VLAN.
  - There are two types of secondary VLAN — community VLAN and isolated VLAN.

PVLAN port types include:

- **Community port** — a port that belongs to a community VLAN and is allowed to communicate with other ports in the same community VLAN and with promiscuous ports.
- **Host port** — in the context of a private VLAN, is a port in a secondary VLAN:
  - The port must first be assigned that role in INTERFACE mode.
  - A port assigned the host role cannot be added to a regular VLAN.
- **Isolated port** — a port that, in Layer 2, can only communicate with promiscuous ports that are in the same PVLAN.
- **Promiscuous port** — a port that is allowed to communicate with any other port type in the PVLAN:
  - A promiscuous port can be part of more than one primary VLAN.
  - A promiscuous port cannot be added to a regular VLAN.
- **Trunk port** — carries traffic between switches:
  - A trunk port in a PVLAN is always tagged.
  - In tagged mode, the trunk port carries the primary or secondary VLAN traffic. The tag on the packet helps identify the VLAN to which the packet belongs.
  - A trunk port can also belong to a regular VLAN (non-private VLAN).

Each of the port types can be any type of physical Ethernet port, including port channels (LAGs). For more information about port channels, refer to [Port Channel Interfaces](#) in the [Interfaces](#) chapter.

For an introduction to VLANs, refer to [Layer 2](#).

## Using the Private VLAN Commands

To use the PVLAN feature, use the following commands.

- Enable/disable Layer 3 communication between secondary VLANs.

INTERFACE VLAN mode

```
[no] ip local-proxy-arp
```



**NOTE:** Even after you disable `ip-local-proxy-arp` (`no ip-local-proxy-arp`) in a secondary VLAN, Layer 3 communication may happen between some secondary VLAN hosts, until the address resolution protocol (ARP) timeout happens on those secondary VLAN hosts.

- Set the mode of the selected VLAN to community, isolated, or primary.

INTERFACE VLAN mode

- ```
[no] private-vlan mode {community | isolated | primary}
```
- Map secondary VLANs to the selected primary VLAN.
INTERFACE VLAN mode
- ```
[no] private-vlan mapping secondary-vlan vlan-list
```
- Display type and status of PVLAN interfaces.  
EXEC mode or EXEC Privilege mode
- ```
show interfaces private-vlan [interface interface]
```
- Display PVLANS and/or interfaces that are part of a PVLAN.
EXEC mode or EXEC Privilege mode
- ```
show vlan private-vlan [community | interface | isolated | primary |
primary_vlan | interface interface]
```
- Display primary-secondary VLAN mapping.  
EXEC mode or EXEC Privilege mode
- ```
show vlan private-vlan mapping
```
- Set the PVLAN mode of the selected port.
INTERFACE
- ```
switchport mode private-vlan {host | promiscuous | trunk}
```
- ✎ **NOTE:** Secondary VLANs are Layer 2 VLANs, so even if they are operationally down while primary VLANs are operationally up, Layer 3 traffic is still transmitted across secondary VLANs.
  - ✎ **NOTE:** The outputs of the `show arp` and `show vlan` commands are augmented in the Dell Networking OS version 7.8.1.0 to provide PVLAN data. For more information, refer to the *Dell Networking OS Command Line Reference Guide*.

## Configuration Task List

The following sections contain the procedures that configure a private VLAN.

- [Creating PVLAN ports](#)
- [Creating a Primary VLAN](#)
- [Creating a Community VLAN](#)
- [Creating an Isolated VLAN](#)

### Creating PVLAN ports

PVLAN ports are those that will be assigned to the PVLAN.

1. Access INTERFACE mode for the port that you want to assign to a PVLAN.  
CONFIGURATION mode  
  
`interface interface`
2. Enable the port.  
INTERFACE mode  
  
`no shutdown`

3. Set the port in Layer 2 mode.

INTERFACE mode

```
switchport
```

4. Select the PVLAN mode.


INTERFACE mode

```
switchport mode private-vlan {host | promiscuous | trunk}
```

- `host` (isolated or community VLAN port)
- `promiscuous` (intra-VLAN communication port)
- `trunk` (inter-switch PVLAN hub port)

### Example of the `switchport mode private-vlan` Command

For interface details, refer to [Enabling a Physical Interface](#) in the [Interfaces](#) chapter.

 **NOTE:** You cannot add interfaces that are configured as PVLAN host or promiscuous ports to regular VLANs. Conversely, you cannot add “regular” ports (ports not configured as PVLAN ports) to PVLANs as a host or promiscuous member.

The example below shows the `switchport mode private-vlan` command on a port and on a port channel.

```
Dell#conf
Dell(conf)#interface TenGigabitEthernet 2/1
Dell(conf-if-te-2/1)#switchport mode private-vlan promiscuous

Dell(conf)#interface TenGigabitEthernet 2/2
Dell(conf-if-te-2/2)#switchport mode private-vlan host

Dell(conf)#interface TenGigabitEthernet 2/3
Dell(conf-if-te-2/3)#switchport mode private-vlan trunk

Dell(conf)#interface TenGigabitEthernet 2/2
Dell(conf-if-te-2/2)#switchport mode private-vlan host
```

### Creating a Primary VLAN

A primary VLAN is a port-based VLAN that is specifically enabled as a primary VLAN to contain the promiscuous ports and PVLAN trunk ports for the private VLAN.

A primary VLAN also contains a mapping to secondary VLANs, which are comprised of community VLANs and isolated VLANs.

1. Access INTERFACE VLAN mode for the VLAN to which you want to assign the PVLAN interfaces.

CONFIGURATION mode

```
interface vlan vlan-id
```

2. Enable the VLAN.

INTERFACE VLAN mode

```
no shutdown
```

3. Set the PVLAN mode of the selected VLAN to primary.

INTERFACE VLAN mode



```
private-vlan mode primary
```

4. Map secondary VLANs to the selected primary VLAN.

```
INTERFACE VLAN mode
```

```
private-vlan mapping secondary-vlan vlan-list
```

The list of secondary VLANs can be:

- Specified in comma-delimited (*VLAN-ID, VLAN-ID*) or hyphenated-range format (*VLAN-ID-VLAN-ID*).
  - Specified with this command even before they have been created.
  - Amended by specifying the new secondary VLAN to be added to the list.
5. Add promiscuous ports as tagged or untagged interfaces.

```
INTERFACE VLAN mode
```

```
tagged interface Or untagged interface
```

Add PVLAN trunk ports to the VLAN only as tagged interfaces.

You can enter interfaces singly or in range format, either comma-delimited (*slot/port, port, port*) or hyphenated (*slot/port-port*).

You can only add promiscuous ports or PVLAN trunk ports to the PVLAN (no host or regular ports).

6. (OPTIONAL) Assign an IP address to the VLAN.


```
INTERFACE VLAN mode
```

```
ip address ip address
```

7. (OPTIONAL) Enable/disable Layer 3 communication between secondary VLANs.

```
INTERFACE VLAN mode
```

```
ip local-proxy-arp
```

-  **NOTE:** If a promiscuous or host port is untagged in a VLAN and it receives a tagged packet in the same VLAN, the packet is NOT dropped.

## Creating a Community VLAN

A community VLAN is a secondary VLAN of the primary VLAN in a private VLAN.

The ports in a community VLAN can talk to each other and with the promiscuous ports in the primary VLAN.

1. Access INTERFACE VLAN mode for the VLAN that you want to make a community VLAN.

```
CONFIGURATION mode
```

```
interface vlan vlan-id
```

2. Enable the VLAN.

```
INTERFACE VLAN mode
```

```
no shutdown
```

3. Set the PVLAN mode of the selected VLAN to community.

```
INTERFACE VLAN mode
```

```
private-vlan mode community
```

4. Add one or more host ports to the VLAN.

```
INTERFACE VLAN mode
```

```
tagged interface OR untagged interface
```

You can enter the interfaces singly or in range format, either comma-delimited (*slot/port,port,port*) or hyphenated (*slot/ port-port*).

You can only add host (isolated) ports to the VLAN.

### Creating an Isolated VLAN

An isolated VLAN is a secondary VLAN of a primary VLAN.

An isolated VLAN port can only talk with the promiscuous ports in that primary VLAN.

1. Access INTERFACE VLAN mode for the VLAN that you want to make an isolated VLAN.

```
CONFIGURATION mode
```

```
interface vlan vlan-id
```

2. Enable the VLAN.

```
INTERFACE VLAN mode
```

```
no shutdown
```

3. Set the PVLAN mode of the selected VLAN to isolated.

```
INTERFACE VLAN mode
```

```
private-vlan mode isolated
```

4. Add one or more host ports to the VLAN.

```
INTERFACE VLAN mode
```

```
tagged interface OR untagged interface
```

You can enter the interfaces singly or in range format, either comma-delimited (*slot/port,port,port*) or hyphenated (*slot/ port-port*).

You can only add ports defined as host to the VLAN.

### Example of Configuring Private VLAN Members

The following example shows the use of the PVLAN commands that are used in VLAN INTERFACE mode to configure the PVLAN member VLANs (primary, community, and isolated VLANs).

```
Dell#conf
Dell(conf)# interface vlan 10
Dell(conf-vlan-10)# private-vlan mode primary
Dell(conf-vlan-10)# private-vlan mapping secondary-vlan 100-101
Dell(conf-vlan-10)# untagged TenGig 2/1
Dell(conf-vlan-10)# tagged TenGig 2/3

Dell(conf)# interface vlan 101
Dell(conf-vlan-101)# private-vlan mode community
Dell(conf-vlan-101)# untagged TenGig 2/10
```

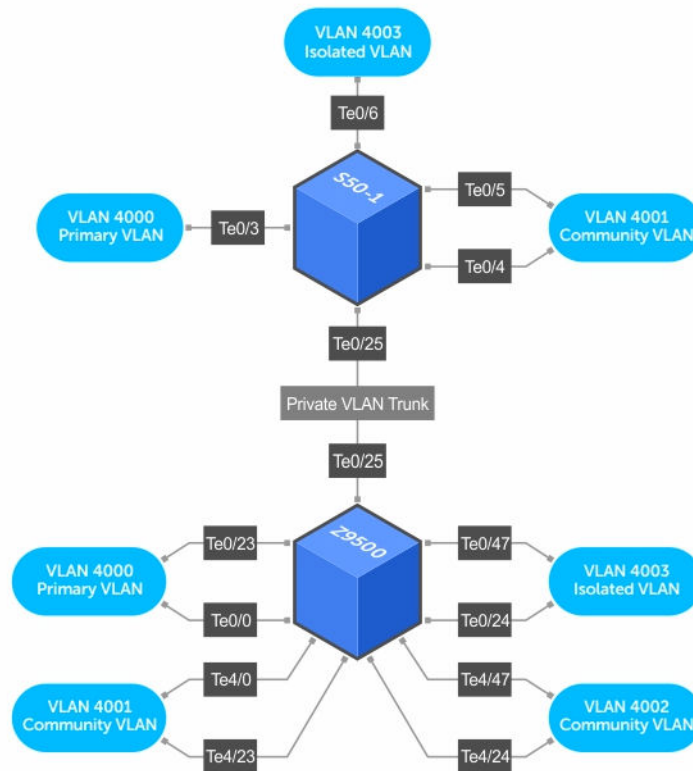
```

Dell(conf)# interface vlan 100
Dell(conf-vlan-100)# private-vlan mode isolated
Dell(conf-vlan-100)# untagged Te 2/2

```

## Private VLAN Configuration Example

The following example shows a private VLAN topology.




**Figure 99. Sample Private VLAN Topology**

The following configuration is based on the example diagram for the MXL switch:

- TenGig 0/0 and TenGig 0/23 are configured as promiscuous ports, assigned to the primary VLAN, VLAN 4000.
- TenGig 0/25 is configured as a PVLAN trunk port, also assigned to the primary VLAN 4000.
- TenGig 0/24 and TenGig 0/47 are configured as host ports and assigned to the isolated VLAN, VLAN 4003.
- TenGig 4/0 and TenGig 0/23 are configured as host ports and assigned to the community VLAN, VLAN 4001.
- TenGig 4/24 and TenGig 4/47 are configured as host ports and assigned to community VLAN 4002.

The result is that:

- The ports in community VLAN 4001 can communicate directly with each other and with promiscuous ports.
- The ports in community VLAN 4002 can communicate directly with each other and with promiscuous ports.
- The ports in isolated VLAN 4003 can only communicate with the promiscuous ports in the primary VLAN 4000.
- All the ports in the secondary VLANs (both community and isolated VLANs) can only communicate with ports in the other secondary VLANs of that PVLAN over Layer 3, and only when the `ip local-proxy-arp` command is invoked in the primary VLAN.

 **NOTE:** Even after you disable `ip-local-proxy-arp` (`no ip-local-proxy-arp`) in a secondary VLAN, Layer 3 communication may happen between some secondary VLAN hosts, until the ARP timeout happens on those secondary VLAN hosts.

### Inspecting the Private VLAN Configuration

The standard methods of inspecting configurations also apply in PVLANS. To inspect your PVLAN configurations, use the following commands.

- Display the specific interface configuration.  
INTERFACE mode and INTERFACE VLAN mode

```
show config
```

- Inspect the running-config, and, with the `grep pipe` option, display a specific part of the running-config.

```
show running-config | grep string
```

The following example shows the PVLAN parts of the running-config from the S50V switch in the topology diagram previously shown.

- Display the type and status of the configured PVLAN interfaces.  
`show interfaces private-vlan [interface interface]`

This command is specific to the PVLAN feature.

For more information, refer to the *Security* chapter in the *Dell Networking OS Command Line Reference Guide*.

- Display the configured PVLANS or interfaces that are part of a PVLAN.

```
show vlan private-vlan [community | interface | isolated | primary | primary_vlan | interface interface]
```

This command is specific to the PVLAN feature.

The following examples show the results of using this command without the command options on the MXL switch in the topology diagram previously shown.

- Display the primary-secondary VLAN mapping. The following example shows the output from the MXL switch.

```
show vlan private-vlan mapping
```

This command is specific to the PVLAN feature.

## Example of Viewing a Private VLAN

### Example of the show vlan private-vlan mapping Command

### Example of Viewing VLAN Status


### Example of Viewing Private VLAN Configuration

The show arp and show vlan commands are revised to display PVLAN data.

```
Dell#show vlan private-vlan
Primary Secondary Type Active Ports

20 30 Primary Yes Te 1/1,5
 40 Community Yes Te 1/2
 40 Isolated Yes Te 1/3
Dell#
```

```
S50-1#show vlan private-vlan mapping
Private Vlan:
Primary : 4000
Isolated : 4003
Community : 4001
```

 **NOTE:** In the following example, notice the addition of the PVLAN codes – P, I, and C – in the left column.

```
Dell#show vlan
Codes: * - Default VLAN, G - GVRP VLANs, R - Remote Port Mirroring VLANs, P -
Primary, C -
Community, I - Isolated
Q: U - Untagged, T - Tagged
 x - Dot1x untagged, X - Dot1x tagged
 G - GVRP tagged, M - Vlan-stack, H - VSN tagged
 i - Internal untagged, I - Internal tagged, v - VLT untagged, V - VLT tagged

NUM Status Description Q Ports
* 1 Active U Te 5/41
P 20 Active T Te 1/1,5
C 30 Active T Te 1/2
I 40 Active T Te 1/3
Dell#
```

```
Dell#show vlan
!
interface TenGigabitEthernet 1/1
 no ip address
 switchport
 switchport mode private-vlan promiscuous
 no keepalive
 no shutdown
!
interface TenGigabitEthernet 1/2
 no ip address
 switchport
 switchport mode private-vlan host
 no shutdown
!
interface TenGigabitEthernet 1/3
 no ip address
 switchport
```

```
 switchport mode private-vlan host
 no shutdown
!
interface TenGigabitEthernet 1/5
 no ip address
 switchport
 switchport mode private-vlan trunk
 no shutdown
interface Vlan 20
 private-vlan mode primary
 private-vlan mapping secondary-vlan 30,40
 no ip address
 tagged TenGigabitEthernet 1/1,5
 shutdown
!
interface Vlan 30
 private-vlan mode community
 no ip address
 tagged TenGigabitEthernet 1/2
 no shutdown
!
```

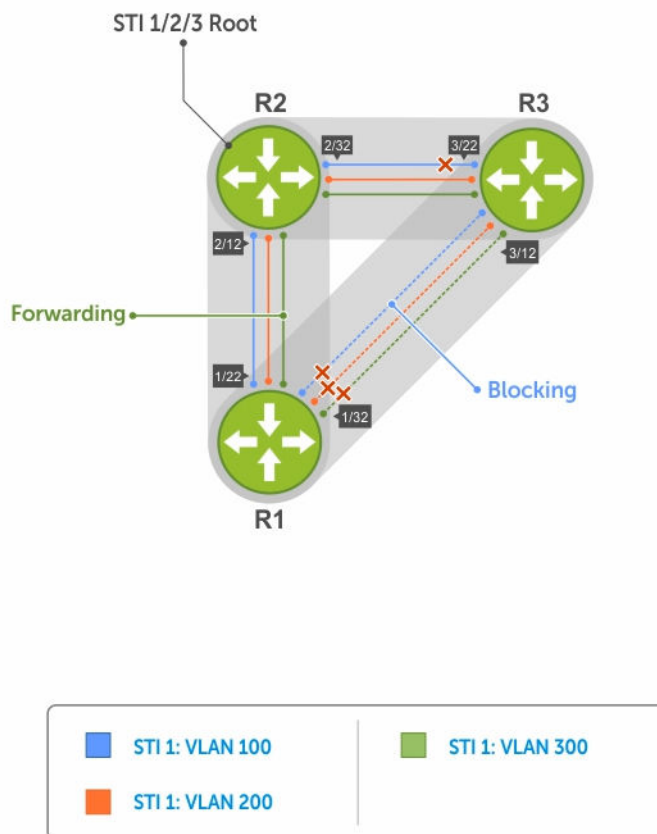
# Per-VLAN Spanning Tree Plus (PVST+)

Per-VLAN spanning tree plus (PVST+) is supported on the MXL switch platform.

## Protocol Overview

PVST+ is a variation of spanning tree — developed by a third party — that allows you to configure a separate spanning tree instance for each virtual local area network (VLAN).

For more information about spanning tree, refer to the [Spanning Tree Protocol \(STP\)](#) chapter.



**Figure 100. Per-VLAN Spanning Tree**

The Dell Networking operating system (OS) supports three other variations of spanning tree, as shown in the following table.

**Table 46. Spanning Tree Variations Dell Networking OS Supports**

| Dell Networking Term                                   | IEEE Specification |
|--------------------------------------------------------|--------------------|
| <a href="#">Spanning Tree Protocol (STP)</a>           | 802 .1d            |
| <a href="#">Rapid Spanning Tree Protocol (RSTP)</a>    | 802 .1w            |
| <a href="#">Multiple Spanning Tree Protocol (MSTP)</a> | 802 .1s            |
| <a href="#">Per-VLAN Spanning Tree Plus (PVST+)</a>    | Third Party        |

## Implementation Information

- The Dell Networking OS implementation of PVST+ is based on IEEE Standard 802.1w.
- The Dell Networking OS implementation of PVST+ uses IEEE 802.1s costs as the default costs (as shown in the following table). Other implementations use IEEE 802.1w costs as the default costs. If you are using Dell Networking systems in a multivendor network, verify that the costs are values you intended.
- You can enable PVST+ on 254 VLANs.

## Configure Per-VLAN Spanning Tree Plus

Configuring PVST+ is a four-step process.

1. Configure interfaces for Layer 2.
2. Place the interfaces in VLANs.
3. Enable PVST+.
4. Optionally, for load balancing, select a nondefault bridge-priority for a VLAN.

### Related Configuration Tasks

- [Modifying Global PVST+ Parameters](#)
- [Enable BPDU Filtering Globally](#)
- [Configuring an EdgePort](#)
- [Flush MAC Addresses after a Topology Change](#)
- [Prevent Network Disruptions with BPDU Guard](#)
- [SNMP Traps for Root Elections and Topology Changes](#)
- [PVST+ in Multi-Vendor Networks](#)
- [Enabling PVST+ Extend System ID](#)
- [PVST+ Sample Configurations](#)

## Enabling PVST+

When you enable PVST+, the Dell Networking OS instantiates STP on each active VLAN.

1. Enter PVST context.  
PROTOCOL PVST mode



- ```
protocol spanning-tree pvst
```
2. Enable PVST+.
PROTOCOL PVST mode

no disable

Disabling PVST+

To disable PVST+ globally or on an interface, use the following commands.

- Disable PVST+ globally.
PROTOCOL PVST mode

disable
- Disable PVST+ on an interface, or remove a PVST+ parameter configuration.
INTERFACE mode

no spanning-tree pvst

Example of Viewing PVST+ Configuration

To display your PVST+ configuration, use the `show config` command from PROTOCOL PVST mode.

```
Dell(conf-pvst)#show config verbose
!
protocol spanning-tree pvst
  no disable
  vlan 100 bridge-priority 4096
```

Influencing PVST+ Root Selection

As shown in the previous per-VLAN spanning tree illustration, all VLANs use the same forwarding topology because R2 is elected the root, and all TenGigabitEthernet ports have the same cost. The following per-VLAN spanning tree illustration changes the bridge priority of each bridge so that a different forwarding topology is generated for each VLAN. This behavior demonstrates how you can use PVST+ to achieve load balancing.




Figure 101. Load Balancing with PVST+

The bridge with the bridge value for bridge priority is elected root. Because all bridges use the default priority (until configured otherwise), the lowest MAC address is used as a tie-breaker. To increase the likelihood that a bridge is selected as the STP root, assign bridges a low non-default value for bridge priority.

To assign a bridge priority, use the following command.

- Assign a bridge priority.
PROTOCOL PVST mode

```
vlan bridge-priority
```

The range is from 0 to 61440.

The default is **32768**.

Example of the show spanning-tree pvst vlan Command

To display the PVST+ forwarding topology, use the `show spanning-tree pvst [vlan vlan-id]` command from EXEC Privilege mode.

```
Dell(conf-if-te-5/41)#do show spanning-tree pvst vlan 2
VLAN 2
```

```
Root Identifier has priority 32768, Address 001e.c9f1.00f3
Root Bridge hello time 2, max age 20, forward delay 15
Bridge Identifier has priority 32768, Address 001e.c9f1.00f3
Configured hello time 2, max age 20, forward delay 15
Bpdu filter disabled globally
We are the root of VLAN 2
Current root has priority 32768, Address 001e.c9f1.00f3
Number of topology changes 2, last change occurred 00:14:39 ago on Po 23
```

```
Port 24 (Port-channel 23) is designated Forwarding
Port path cost 1600, Port priority 128, Port Identifier 128.24
Designated root has priority 32768, address 001e.c9f1.00:f3
Designated bridge has priority 32768, address 001e.c9f1.00:f3
Designated port id is 128.24 , designated path cost 0
Number of transitions to forwarding state 1
BPDU sent 449, received 0
The port is not in the Edge port mode, bpdu filter is disabled
```

```
Port 450 (TenGigabitEthernet 5/41) is disabled Discarding
Port path cost 2000, Port priority 128, Port Identifier 128.450
Designated root has priority 32768, address 001e.c9f1.00:f3
Designated bridge has priority 32768, address 001e.c9f1.00:f3
Designated port id is 128.450 , designated path cost 0
Number of transitions to forwarding state 0
BPDU sent 0, received 0
The port is not in the Edge port mode, bpdu filter is disabled
```

```
Port 459 (TenGigabitEthernet 5/50) is designated Forwarding
Port path cost 2000, Port priority 128, Port Identifier 128.459
Designated root has priority 32768, address 001e.c9f1.00:f3
Designated bridge has priority 32768, address 001e.c9f1.00:f3
Designated port id is 128.459 , designated path cost 0
Number of transitions to forwarding state 1
BPDU sent 457, received 0
The port is not in the Edge port mode, bpdu filter is disabled
```

Modifying Global PVST+ Parameters

The root bridge sets the values for forward-delay and hello-time, and overwrites the values set on other PVST+ bridges.

- **Forward-delay** — the amount of time an interface waits in the Listening state and the Learning state before it transitions to the Forwarding state.
- **Hello-time** — the time interval in which the bridge sends bridge protocol data units (BPDUs).
- **Max-age** — the length of time the bridge maintains configuration information before it refreshes that information by recomputing the PVST+ topology.

To change PVST+ parameters on the root bridge, use the following commands.

- Change the forward-delay parameter.
PROTOCOL PVST mode


```
vlan forward-delay
```

The range is from 4 to 30.

The default is **15 seconds**.

- Change the hello-time parameter.
PROTOCOL PVST mode

```
vlan hello-time
```

 **NOTE:** With large configurations (especially those configurations with more ports), Dell Networking recommends increasing the hello-time.

The range is from 1 to 10.

The default is **2 seconds**.

- Change the max-age parameter.
PROTOCOL PVST mode

```
vlan max-age
```

The range is from 6 to 40.

The default is **20 seconds**.

The values for global PVST+ parameters are given in the output of the `show spanning-tree pvst` command.

Modifying Interface PVST+ Parameters

You can adjust two interface parameters (port cost and port priority) to increase or decrease the probability that a port becomes a forwarding port.


- **Port cost** — a value that is based on the interface type. The greater the port cost, the less likely the port is selected to be a forwarding port.
- **Port priority** — influences the likelihood that a port is selected to be a forwarding port in case that several ports have the same port cost.

The following tables lists the default values for port cost by interface.

Table 47. Default Values for Port Cost

Port Cost	Default Value
1000-Mb/s Ethernet interfaces	20000
10-Gigabit Ethernet interfaces	2000
40-Gigabit Ethernet interfaces	1400
Port Channel with one 10-Gigabit Ethernet interface	2000
Port Channel with one 40-Gigabit Ethernet interface	1400
Port Channel with two 10-Gigabit Ethernet interfaces	1800

Port Cost	Default Value
Port Channel with two 40-Gigabit Ethernet interfaces	600

 **NOTE:** The Dell Networking OS implementation of PVST+ uses IEEE 802.1s costs as the default costs. Other implementations use IEEE 802.1w costs as the default costs. If you are using Dell Networking systems in a multi-vendor network, verify that the costs are values you intended.

To change the port cost or port priority of an interface, use the following commands.

- Change the port cost of an interface.

INTERFACE mode

```
spanning-tree pvst vlan cost.
```

The range is from 0 to 200000.

Refer to the table for the default values.

- Change the port priority of an interface.

INTERFACE mode

```
spanning-tree pvst vlan priority.
```

The range is from 0 to 240, in increments of 16.


The default is **128**.

The values for interface PVST+ parameters are given in the output of the `show spanning-tree pvst` command, as previously shown.

Configuring an EdgePort

The EdgePort feature enables interfaces to begin forwarding traffic approximately 30 seconds sooner. In this mode an interface forwards frames by default until it receives a BPDU that indicates that it should behave otherwise; it does not go through the Learning and Listening states. The `bpduguard shutdown-on-violation` option causes the interface hardware to be shut down when it receives a BPDU. When you only implement `bpduguard`, although the interface is placed in an Error Disabled state when receiving the BPDU, the physical interface remains up and spanning-tree drops packets in the hardware after a BPDU violation. BPDUs are dropped in the software after receiving the BPDU violation.

This feature is the same as PortFast mode in spanning tree.

 **CAUTION: Configure EdgePort only on links connecting to an end station. EdgePort can cause loops if you enable it on an interface connected to a network.**

To enable EdgePort on an interface, use the following command.

- Enable EdgePort on an interface.

INTERFACE mode

```
spanning-tree pvst edge-port [bpduguard | shutdown-on-violation]
```

The EdgePort status of each interface is given in the output of the `show spanning-tree pvst` command, as previously shown.

Dell Networking OS Behavior: Regarding the `bpduguard shutdown-on-violation` command behavior:

- If the interface to be shut down is a port channel, all the member ports are disabled in the hardware.
- When you add a physical port to a port channel already in an Error Disable state, the new member port is also disabled in the hardware.
- When you remove a physical port from a port channel in an Error Disable state, the Error Disabled state is cleared on this physical port (the physical port is enabled in the hardware).
- You can clear the Error Disabled state with any of the following methods:
 - Perform a `shutdown` command on the interface.
 - Disable the `shutdown-on-violation` command on the interface (the `no spanning-tree stp-id portfast [bpduguard | [shutdown-on-violation]]` command).
 - Disable spanning tree on the interface (the `no spanning-tree` command in INTERFACE mode).
 - Disabling global spanning tree (the `no spanning-tree` command in CONFIGURATION mode).

PVST+ in Multi-Vendor Networks

Some non-Dell Networking systems which have hybrid ports participating in PVST+ transmit two kinds of BPDUs: an 802.1D BPDU and an untagged PVST+ BPDU.

Dell Networking systems do not expect PVST+ BPDU (tagged or untagged) on an untagged port. If this situation occurs, the Dell Networking OS places the port in an Error-Disable state. This behavior might result in the network not converging. To prevent the system from executing this action, use the `no spanning-tree pvst err-disable cause invalid-pvst-bpdu` command. After you configure this command, if the port receives a PVST+ BPDU, the BPDU is dropped and the port remains operational.

Enabling PVST+ Extend System ID

In the following example, ports P1 and P2 are untagged members of different VLANs. These ports are untagged because the hub is VLAN unaware. There is no data loop in this scenario; however, you can employ PVST+ to avoid potential misconfigurations.

If you enable PVST+ on the Dell Networking switch in this network, P1 and P2 receive BPDUs from each other. Ordinarily, the Bridge ID in the frame matches the Root ID, a loop is detected, and the rules of convergence require that P2 move to blocking state because it has the lowest port ID.

To keep both ports in a Forwarding state, use extend system ID. Extend system ID augments the bridge ID with a VLAN ID to differentiate BPDUs on each VLAN so that PVST+ does not detect a loop and both ports can remain in a Forwarding state.




Figure 102. PVST+ with Extend System ID

- Augment the bridge ID with the VLAN ID.
PROTOCOL PVST mode

```
extend system-id
```

Example of Viewing the Extend System ID in a PVST+ Configuration

```
Dell(conf-pvst)#do show spanning-tree pvst vlan 5 brief
```

```
VLAN 5
Executing IEEE compatible Spanning Tree Protocol
Root ID Priority 32773, Address 0001.e832.73f7
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID Priority 32773 (priority 32768 sys-id-ext 5), Address 0001.e832.73f7
We are the root of Vlan 5
Configured hello time 2, max age 20, forward delay 15
```

PVST+ Sample Configurations

The following examples provide the running configurations for the topology shown in the previous illustration.

Example of PVST+ Configuration (R1)

Example of PVST+ Configuration (R2)

Example of PVST+ Configuration (R3)

```
interface TenGigabitEthernet 1/22
  no ip address
  switchport
  no shutdown
!
```

```

interface TenGigabitEthernet 1/32
  no ip address
  switchport
  no shutdown
!
  protocol spanning-tree pvst
  no disable
  vlan 100 bridge-priority 4096
interface Vlan 100
  no ip address
  tagged TenGigabitEthernet 1/22,32
  no shutdown
!
  interface Vlan 200
  no ip address
  tagged TenGigabitEthernet 1/22,32
  no shutdown
!
  interface Vlan 300
  no ip address
  tagged TenGigabitEthernet 1/22,32
  no shutdown
!
protocol spanning-tree pvst
  no disable
  vlan 100 bridge-priority 4096

interface TenGigabitEthernet 2/12
  no ip address
  switchport
  no shutdown
!
interface TenGigabitEthernet 2/32
  no ip address
  switchport
  no shutdown
!
interface Vlan 100
  no ip address
  tagged TenGigabitEthernet 2/12,32
  no shutdown
!
interface Vlan 200
  no ip address
  tagged TenGigabitEthernet 2/12,32
  no shutdown
!
interface Vlan 300
  no ip address
  tagged TenGigabitEthernet 2/12,32
  no shutdown
!
protocol spanning-tree pvst
  no disable
  vlan 200 bridge-priority 4096

interface TenGigabitEthernet 3/12
  no ip address
  switchport
  no shutdown
!
interface TenGigabitEthernet 3/22

```



```
no ip address
switchport
no shutdown
!
interface Vlan 100
no ip address
tagged TenGigabitEthernet 3/12,22
no shutdown
!
interface Vlan 200
no ip address
tagged TenGigabitEthernet 3/12,22
no shutdown
!
interface Vlan 300
no ip address
tagged TenGigabitEthernet 3/12,22
no shutdown
!
protocol spanning-tree pvst
no disable
vlan 300 bridge-priority 4096
```

Enable BPDU Filtering globally

The enabling of BPDU Filtering stops transmitting of BPDUs on the operational port fast enabled ports by default.

When BPDUs are received, the spanning tree is automatically prepared. By default global bpdu filtering is disabled.

Enable BPDU Filter globally to filter transmission of BPDU port fast enabled interfaces.
PROTOCOL PVST mode

edge-port bpdu filter default

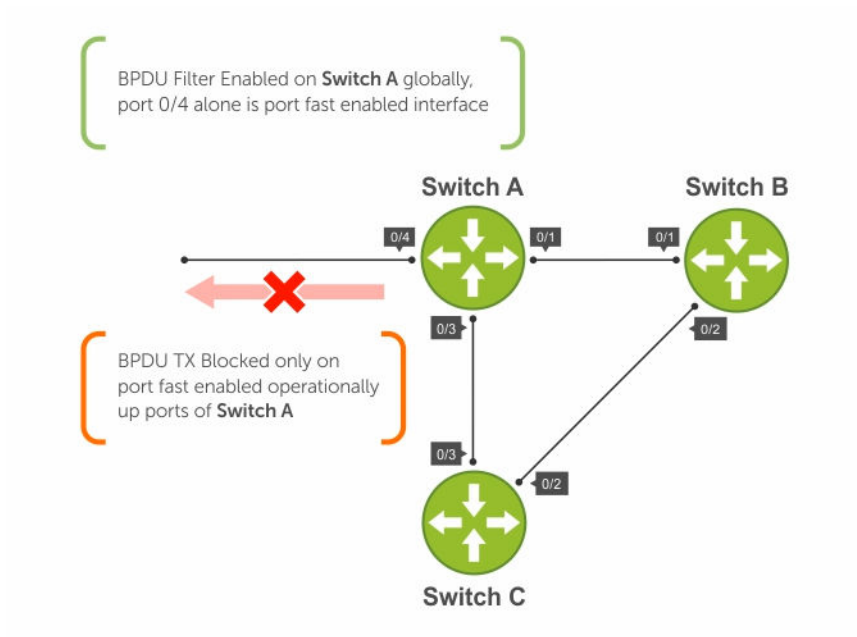


Figure 103. BPDU Filtering Enabled Globally

Quality of Service (QoS)

Quality of service (QoS) is supported on the MXL switch platform.

Differentiated service is accomplished by classifying and queuing traffic, and assigning priorities to those queues.

The MXL switch traffic has four data queues per port. All queues are serviced using the Weighted Round Robin scheduling algorithm. You can only manage prioritize queuing on egress.


 **NOTE:** When you enable DCB, the egress QoS features in the output QoS policy-map (such as service-class bandwidth-percentage and bandwidth-percentage), the default bandwidth allocation ratio for egress queues and strict-priority may not work as intended. This is to provide compatibility with DCBX. Hence, it is recommended to have the DCB disabled when you wish to apply these features exclusively.

Table 48. Dell Networking OS Support for Port-Based, Policy-Based, and Multicast QoS Features

Feature	Direction
Port-Based QoS Configurations	Ingress + Egress
Set dot1p Priorities for Incoming Traffic	Ingress
Honor dot1p Priorities on Ingress Traffic	Ingress
Configure Port-based Rate Policing	Ingress
Configure Port-based Rate Shaping	Egress
Policy-Based QoS Configurations	Ingress + Egress
Classify Traffic	Ingress
Create a Layer 3 Class Map	Ingress
Set DSCP Values for Egress Packets Based on Flow	Ingress
Create a Layer 2 Class Map	Ingress
Create a QoS Policy	Ingress + Egress
Create an Input QoS Policy	Ingress
Configure Policy-Based Rate Policing	Ingress
Set a DSCP Value for Egress Packets	Ingress
Set a dot1p Value for Egress Packets	Ingress
Create an Output QoS Policy	Egress
Configure Policy-Based Rate Shaping	Egress
Allocate Bandwidth to the Queue	Egress
Configure a Scheduler to Queue	Egress
Specify WRED Drop Precedence	Egress

Feature	Direction
Create Policy Maps	Ingress + Egress
Create Input Policy Maps	Ingress
Honor DSCP Values on Ingress Packets	Ingress
Honoring dot1p Values on Ingress Packets	Ingress
Create Output Policy Maps	Egress
Specify an Aggregate QoS Policy	Egress
QoS Rate Adjustment	
Strict-Priority Queueing	
Weighted Random Early Detection	Egress
Create WRED Profiles	Egress



Figure 104. Dell Networking QoS Architecture

Implementation Information

The Dell Networking QoS implementation complies with IEEE 802.1p *User Priority Bits for QoS Indication*.


It also implements these Internet Engineering Task Force (IETF) documents:

- RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 Headers*
- RFC 2475, *An Architecture for Differentiated Services*
- RFC 2597, *Assured Forwarding PHB Group*
- RFC 2598, *An Expedited Forwarding PHB*

You cannot configure port-based and policy-based QoS on the same interface.

Port-Based QoS Configurations

You can configure the following QoS features on an interface.

 **NOTE:** You cannot simultaneously use egress rate shaping and ingress rate policing on the same virtual local area network (VLAN).

- [Setting dot1p Priorities for Incoming Traffic](#)
- [Configuring Port-Based Rate Policing](#)
- [Configuring Port-Based Rate Shaping](#)

Setting dot1p Priorities for Incoming Traffic

Change the priority of incoming traffic on the interface using the dot1p-priority command from INTERFACE mode. The Dell Networking OS places traffic marked with a priority in a queue based on the following table.

If you set a dot1p priority for a port-channel, all port-channel members are configured with the same value. You cannot assign a dot1p value to an individual interface in a port-channel.


Table 49. dot1p-priority Values and Queue Numbers

dot1p	Queue Number
0	0
1	0
2	0
3	1
4	2
5	3
6	3

dot1p	Queue Number
7	3

- Change the priority of incoming traffic on the interface.
`dot1p-priority`

Example of Configuring a dot1p Priority on an Interface


 **NOTE:** The `dot1p-priority` command marks all incoming traffic on an interface with a specified dot1p priority and maps all incoming traffic to the corresponding queue. When you enable PFC and/or ETS on an interface, incoming traffic with a specified dot1p priority can be distributed across different queues. Therefore, when you use PFC and ETS to manage data center traffic, it is not recommended that you use the `dot1p-priority` command to set a queue assignment. For more information, refer to [Data Center Bridging \(DCB\)](#).

```
Dell#config
Dell(conf)#interface tengigabitethernet 1/0
Dell(conf-if)#switchport
Dell(conf-if)#dot1p-priority 1
Dell(conf-if)#end
Dell#
```

Honoring dot1p Priorities on Ingress Traffic

By default, the Dell Networking OS does not honor dot1p priorities on ingress traffic. You can configure this feature on physical interfaces and port-channels, but you cannot configure it on individual interfaces in a port channel.

You can configure `service-class dynamic dot1p` from CONFIGURATION mode, which applies the configuration to all interfaces. A CONFIGURATION mode `service-class dynamic dot1p` entry supersedes any INTERFACE entries. For more information, refer to [Mapping dot1p Values to Service Queues](#).

 **NOTE:** You cannot configure `service-policy input` and `service-class dynamic dot1p` on the same interface.

- Honor dot1p priorities on ingress traffic.
INTERFACE mode
`service-class dynamic dot1p`

Example of Configuring an Interface to Honor dot1p Priorities on Ingress Traffic

```
Dell#config t
Dell(conf)#interface tengigabitethernet 1/0
Dell(conf-if)#service-class dynamic dot1p
Dell(conf-if)#end
Dell#
```

Priority-Tagged Frames on the Default VLAN

Priority-tagged frames are 802.1Q tagged frames with VLAN ID 0.

For VLAN classification, these packets are treated as untagged. However, the dot1p value is still honored when you configure `service-class dynamic dot1p` or `trust dot1p`.

When priority-tagged frames ingress an untagged port or hybrid port, the frames are classified to the default VLAN of the port and to a queue according to their dot1p priority if you configure `service-class dynamic dotp` or `trust dot1p`. When priority-tagged frames ingress a tagged port, the frames are dropped because, for a tagged port, the default VLAN is 0.

Dell Networking OS Behavior: Hybrid ports can receive untagged, tagged, and priority tagged frames. The rate metering calculation might be inaccurate for untagged ports because an internal assumption is made that all frames are treated as tagged. Internally, the ASIC adds a 4-bytes tag to received untagged frames. Though these 4-bytes are not part of the untagged frame received on the wire, they are included in the rate metering calculation resulting in metering inaccuracy.

Configuring Port-Based Rate Policing

If the interface is a member of a VLAN, you may specify the VLAN for which ingress packets are policed.

- Rate policing ingress traffic on an interface.

INTERFACE mode

```
rate police
```

Example of the `rate police` Command

```
Dell#config t
Dell(conf)#interface tengigabitethernet 1/0
Dell(conf-if)#rate police 100 40 peak 150 50
Dell(conf-if)#end
Dell#
```

Configuring Port-Based Rate Shaping

Rate shaping buffers, rather than drops, traffic that exceeds the specified rate until the buffer is exhausted. If any stream exceeds the configured bandwidth on a continuous basis, it can consume all of the buffer space that is allocated to the port.

- Apply rate shaping to outgoing traffic on a port.

INTERFACE mode

```
rate shape
```

- Apply rate shaping to a queue.

QoS Policy mode

```
rate-shape
```

Example of rate-shape Command

```
Dell#config
Dell(conf)#interface tengigabitethernet 1/0
Dell(conf-if)#rate shape 500 50
Dell(conf-if)#end
Dell#
```

Guidelines for Configuring ECN for Classifying and Color-Marking Packets

Keep the following points in mind while configuring the marking and mapping of incoming packets using ECN fields in IPv4 headers:

- Currently Dell Networking OS supports matching only the following TCP flags:
 - ACK
 - FIN
 - SYN
 - PSH
 - RST
 - URG

In the existing software, ECE/CWR TCP flag qualifiers are not supported.

- Because this functionality forcibly marks all the packets matching the specific match criteria as 'yellow', Dell Networking OS does not support Policer based coloring and this feature concurrently.
- If single rate two color policer is configured along with this feature, then by default all packets less than PIR would be considered as "Green" But 'Green' packets matching the specific match criteria for which 'color-marking' is configured will be over-written and marked as "Yellow".
- If two rate three color policer is configured along with this feature then,
 - $x < CIR$ – will be marked as "Green"
 - $CIR < x < PIR$ – will be marked as "Yellow"
 - $PIR < x$ – will be marked as "Red"

But 'Green' packets matching the specific match criteria for which 'color-marking' is configured will be over-written and marked as "Yellow".

Sample configuration to mark non-ecn packets as "yellow" with Multiple traffic class

Consider the example where there are no different traffic classes that is all the packets are egressing on the default 'queue0'.

Dell Networking OS can be configured as below to mark the non-ecn packets as yellow packets.

```
!  
ip access-list standard ecn_0
```



```

seq 5 permit any ecn 0

class-map match-any ecn_0_cmap
  match ip access-group ecn_0 set-color yellow

!
policy-map-input ecn_0_pmap
  service-queue 0 class-map ecn_0_cmap

```

Applying this policy-map "ecn_0_pmap" will mark all the packets with 'ecn == 0' as yellow packets on queue0 (default queue).

Classifying Incoming Packets Using ECN and Color-Marking

Explicit Congestion Notification (ECN) is a capability that enhances WRED by marking the packets instead of causing WRED to drop them when the threshold value is exceeded. If you configure ECN for WRED, devices employ this functionality of ECN to mark the packets and reduce the rate of sending packets in a congested, heavily-loaded network.

ECN is a mechanism using which network switches indicate congestion to end hosts for initiating appropriate action. End hosts uses two least significant bits of ToS to indicate that it is ECT. When intermediate network node encounters congestion, remarks ECT to CE for end host to take appropriate action. During congestion, ECN enabled packets are not subject to any kind of drops like WRED except tail drops. Though ECN & WRED are independent technologies, BRCM has made WRED a mandatory for ECN to work.

On ECN deployment, the non-ECN packets that are transmitted on the ECN-WRED enabled interface will be considered as Green packets and will be subject to the early WRED drops. Typically the TCP-acks, OAM, ICMP ping packets will be non-ECN in nature and it is not desirable for this packets getting WRED dropped.

In such a condition, it is necessary that the switch is capable to take differentiated actions for ECN/Non-ECN packets. After classifying packets to ECN/Non-ECN, marking ECN and Non-ECN packets to different color packets is performed.

Policy based ingress QoS involves the following three steps to achieve QoS:

1. Classification of incoming traffic.
2. Specify the differentiated actions for different traffic class.
3. Attach the policy-map to the interface.

Dell Networking OS support different types of match qualifiers to classify the incoming traffic.

Match qualifiers can be directly configured in the class-map command or it can be specified through one or more ACL which in turn specifies the combination of match qualifiers.

Until Release 9.3(0.0), support is available for classifying traffic based on the 6-bit DSCP field of the IPv4 packet.

As a part of this feature, the 2-bit ECN field of the IPv4 packet will also be available to be configured as one of the match qualifier. This way the entire 8-bit ToS field of the IPv4 header shall be used to classify traffic.

The Dell Networking OS Release 9.3(0.0) supports the following QOS actions in the ingress policy based QOS:

1. Rate Policing
2. Queuing
3. Marking

For the L3 Routed packets, the DSCP marking is the only marking action supported in the software. As a part of this feature, the additional marking action to set the "color" of the traffic will be provided.

Until Release 9.3(0.0), the software has the capability to qualify only on the 6-bit DSCP part of the ToS field in IPv4 Header. You can now accept and process incoming packets based on the 2-bit ECN part of the ToS field in addition to the DSCP categorization. The IPv4 ACLs (standard and Extended) are enhanced to add this qualifier. This new keyword 'ecn' is present for all L3 ACL types (TCP/UDP/IP/ICMP) at the level where the 'DSCP' qualifier is positioned in the current ACL commands.

Dell Networking OS supports the capability to contain DSCP and ECN classifiers simultaneously for the same ACL entry.

You can use the ecn keyword with the ip access-list standard, ip access-list extended, seq, and permit commands for standard and extended IPv4 ACLs to match incoming packets with the specified ECN values.

Similar to 'dscp' qualifier in the existing L3 ACL command, the 'ecn' qualifier can be used along with all other supported ACL match qualifiers such as SIP/DIP/TCP/UDP/SRC PORT/DST PORT/ ICMP.

Until Release 9.3(0.0), ACL supports classification based on the below TCP flags:

- ACK
- FIN
- SYN
- PSH
- RST
- URG

You can now use the 'ecn' match qualifier along with the above TCP flag for classification.

The following combination of match qualifiers is acceptable to be configured for the Dell Networking OS software through L3 ACL command:

- Classification based on DSCP only
- Classification based on ECN only
- Classification based on ECN and DSCP concurrently

You can now use the set-color yellow keyword with the match ip access-group command to mark the color of the traffic as 'yellow' would be added in the 'match ip' sequence of the class-map configuration.

By default, all packets are considered as 'green' (without the rate-policer and trust-diffserve configuration) and hence support would be provided to mark the packets as 'yellow' alone will be provided.

By default Dell Networking OS drops all the 'RED' or 'violate' packets.

The following combination of marking actions to be specified match sequence of the class-map command:

- set a new DSCP for the packet
- set the packet color as 'yellow'
- set the packet color as 'yellow' and set a new DSCP for the packet

This marking action to set the color of the packet is allowed only on the 'match-any' logical operator of the class-map.

This marking-action can be configured for all of the below L3 match sequence types:

- match ip access-group
- match ip dscp
- match ip precedence
- match ip vlan

Sample configuration to mark non-ecn packets as "yellow" with single traffic class

Consider the use case where the packet with DSCP value "40" need to be enqueued in queue#2 and packets with DSCP value as 50 need to be enqueued in queue#3. And all the packets with ecn value as '0' must be marked as 'yellow'.

The above requirement can be achieved using either of the two approaches.

The above requirement can be achieved using either of the two approaches.

Approach without explicit ECN match qualifiers for ECN packets:

```
!  
ip access-list standard dscp_50  
  seq 5 permit any dscp 50  
!  
ip access-list standard dscp_40  
  seq 5 permit any dscp 40  
!  
ip access-list standard dscp_50_non_ecn  
  seq 5 permit any dscp 50 ecn 0  
  
!  
ip access-list standard dscp_40_non_ecn  
  seq 5 permit any dscp 40 ecn 0  
!  
class-map match-any class_dscp_40  
  match ip access-group dscp_40_non_ecn set-color yellow  
  match ip access-group dscp_40  
!  
class-map match-any class_dscp_50  
  match ip access-group dscp_50_non_ecn set-color yellow  
  match ip access-group dscp_50  
!  
policy-map-input pmap_dscp_40_50
```

```
service-queue 2 class-map class_dscp_40
service-queue 3 class-map class_dscp_50
```

Approach with explicit ECN match qualifiers for ECN packets:

```
!
ip access-list standard dscp_50_ecn
  seq 5 permit any dscp 50 ecn 1
  seq 10 permit any dscp 50 ecn 2
  seq 15 permit any dscp 50 ecn 3

!
ip access-list standard dscp_40_ecn
  seq 5 permit any dscp 40 ecn 1
  seq 10 permit any dscp 40 ecn 2
  seq 15 permit any dscp 40 ecn 3

!
ip access-list standard dscp_50_non_ecn
  seq 5 permit any dscp 50 ecn 0

!
ip access-list standard dscp_40_non_ecn
  seq 5 permit any dscp 40 ecn 0
!
class-map match-any class_dscp_40
  match ip access-group dscp_40_non_ecn set-color yellow
  match ip access-group dscp_40_ecn
!
class-map match-any class_dscp_50
  match ip access-group dscp_50_non_ecn set-color yellow
  match ip access-group dscp_50_ecn
!
policy-map-input pmap_dscp_40_50
  service-queue 2 class-map class_dscp_40
  service-queue 3 class-map class_dscp_50
```

Policy-Based QoS Configurations

Policy-based QoS configurations consist of the components shown in the following example.



Figure 105. Constructing Policy-Based QoS Configurations

DSCP Color Maps

This section describes how to configure color maps and how to display the color map and color map configuration.

This sections consists of the following topics:

- Creating a DSCP Color Map
- Displaying Color Maps
- Display Color Map Configuration

Creating a DSCP Color Map

You can create a DSCP color map to outline the differentiated services codepoint (DSCP) mappings to the appropriate color mapping (green, yellow, red) for the input traffic. The system uses this information to classify input traffic on an interface based on the DSCP value of each packet and assigns it an initial drop precedence of green, yellow, or red.

The default setting for each DSCP value (0-63) is green (low drop precedence). The DSCP color map allows you to set the number of specific DSCP values to yellow or red. Traffic marked as yellow delivers traffic to the egress interface, which will either transmit or drop the packet based on configured queuing behavior. Traffic marked as red (high drop precedence) is dropped.

Important Points to Remember

- All DSCP values that are not specified as yellow or red are colored green (low drop precedence).
- A DSCP value cannot be in both the yellow and red lists. Setting the red or yellow list with any DSCP value that is already in the other list results in an error and no update to that DSCP list is made.
- Each color map can only have one list of DSCP values for each color; any DSCP values previously listed for that color that are not in the new DSCP list are colored green.
- If you configured a DSCP color map on an interface that does not exist or you delete a DSCP color map that is configured on an interface, that interface uses an all green color policy.

To create a DSCP color map:

1. Create the color-aware map QoS DSCP color map.

CONFIGURATION mode

```
qos dscp-color-map color-map-name
```

2. Create the color aware map profile.

DSCP-COLOR-MAP

```
dscp {yellow | red} {list-dscp-values}
```

3. Apply the map profile to the interface.

CONFIG-INTERFACE mode

```
qos dscp-color-policy color-map-name
```

Example: Create a DSCP Color Map

The following example creates a DSCP color map profile, color-awareness policy, and applies it to interface **te 1/11**.

Create the DSCP color map profile, **bat-enclave-map**, with a **yellow** drop precedence, and set the DSCP values to **9,10,11,13,15,16**

```
Dell(conf)# qos dscp-color-map bat-enclave-map
Dell(conf-dscp-color-map)# dscp yellow 9,10,11,13,15,16
Dell (conf-dscp-color-map)# exit
```

Assign the color map, **bat-enclave-map** to interface .

Displaying DSCP Color Maps

To display DSCP color maps, use the **show qos dscp-color-map** command in EXEC mode.

Examples for Creating a DSCP Color Map

Display all DSCP color maps.

```
Dell# show qos dscp-color-map
Dscp-color-map mapONE
  yellow 4,7
  red 20,30
Dscp-color-map mapTWO
  yellow 16,55
```

Display a specific DSCP color map.

```
Dell# show qos dscp-color-map mapTWO
Dscp-color-map mapTWO
  yellow 16,55
```

Displaying a DSCP Color Policy Configuration

To display the DSCP color policy configuration for one or all interfaces, use the **show qos dscp-color-policy {summary [interface] | detail {interface}}** command in EXEC mode.

summary: Displays summary information about a color policy on one or more interfaces.

detail: Displays detailed color policy information on an interface

interface: Enter the name of the interface that has the color policy configured.

Examples for Displaying a DSCP Color Policy

Display summary information about a color policy for one or more interfaces.

Display summary information about a color policy for a specific interface.

Display detailed information about a color policy for a specific interface

Classify Traffic


Class maps differentiate traffic so that you can apply separate quality of service policies to different types of traffic.

For both class maps, Layer 2 and Layer 3, the Dell Networking OS matches packets against match criteria in the order that you configure them.

Creating a Layer 3 Class Map

A Layer 3 class map differentiates ingress packets based on the DSCP value or IP precedence, and characteristics defined in an IP ACL. You can also use VLAN IDs and VRF IDs to classify the traffic using layer 3 class-maps.

You may specify more than one DSCP and IP precedence value, but only one value must match to trigger a positive match for the class map.

 **NOTE:** IPv6 and IP-any class maps cannot match on ACLs or VLANs.

Use step 1 or step 2 to start creating a Layer 3 class map.

1. Create a match-any class map.

CONFIGURATION mode

```
class-map match-any
```

2. Create a match-all class map.

CONFIGURATION mode

```
class-map match-all
```

3. Specify your match criteria.

CLASS MAP mode

```
match ip
```

After you create a class-map, the Dell Networking OS places you in CLASS MAP mode.

Match-any class maps allow up to five ACLs. Match-all class-maps allow only one ACL.

4. Link the class-map to a queue.

POLICY MAP mode

```
service-queue
```

Example of Creating a Layer 3 Class Map

```
Dell(conf)#ip access-list standard acl1
Dell(conf-std-nacl)#permit 20.0.0.0/8
Dell(conf-std-nacl)#exit
Dell(conf)#ip access-list standard acl2
Dell(conf-std-nacl)#permit 20.1.1.0/24 order 0
Dell(conf-std-nacl)#exit
Dell(conf)#class-map match-all cmap1
Dell(conf-class-map)#match ip access-group acl1
Dell(conf-class-map)#exit
Dell(conf)#class-map match-all cmap2
Dell(conf-class-map)#match ip access-group acl2
Dell(conf-class-map)#exit
Dell(conf)#policy-map-input pmap
Dell(conf-policy-map-in)#service-queue 0 to 3 class-map cmap1
Dell(conf-policy-map-in)#service-queue 1 class-map cmap2
Dell(conf-policy-map-in)#exit
Dell(conf)#interface tengig 1/0
Dell(conf-if-te-1/0)#service-policy input pmap
```

Examples f Creating a Layer 3 IPv6 Class Map

The following example matches IPv6 traffic with a DSCP value of 40.

```
Dell(conf)# class-map match-all test
Dell(conf-class-map)# match ipv6 dscp 40
```

The following example matches IPv4 and IPv6 traffic with a precedence value of 3.

```
Dell(conf)# class-map match-any test1
Dell(conf-class-map)#match ip-any precedence 3
```

Creating a Layer 2 Class Map

All class maps are Layer 3 by default; however, you can create a Layer 2 class map by specifying the `layer2` option with the `class-map` command.

A Layer 2 class map differentiates traffic according to 802.1p value and/or characteristics defined in a MAC ACL.

Use Step 1 or Step 2 to start creating a Layer 2 class map.

1. Create a match-any class map.

CONFIGURATION mode

```
class-map match-any
```

2. Create a match-all class map.

CONFIGURATION mode

```
class-map match-all
```

3. Specify your match criteria.

CLASS MAP mode

```
match mac
```

After you create a class-map, the system places you in CLASS MAP mode.

Match-any class maps allow up to five access-lists. Match-all class-maps allow only one. You can match against only one VLAN ID.

4. Link the class-map to a queue.

POLICY MAP mode

```
service-queue
```

Determining the Order in Which ACLs are Used to Classify Traffic

When you link class-maps to queues using the `service-queue` command, the system matches the class-maps according to queue priority (queue numbers closer to 0 have lower priorities).

For example, as described in the previous example, class-map *cmap2* is matched against ingress packets before *cmap1*.

ACLs *acl1* and *acl2* have overlapping rules because the address range 20.1.1.0/24 is within 20.0.0.0/8. Therefore (without the keyword `order`), packets within the range 20.1.1.0/24 match positive against *cmap1* and are buffered in queue 4, though you intended for these packets to match positive against *cmap2* and be buffered in queue 1.

In cases such as these, where class-maps with overlapping ACL rules are applied to different queues, use the keyword `order`. The Dell Networking OS writes to the CAM ACL rules with lower order numbers (order numbers closer to 0) before rules with higher order numbers so that packets are matched as you intended.

- Specify the order in which you want to apply ACL rules using the keyword `order`.

The order can range from 0 to 254.

By default, all ACL rules have an order of **255**.

Setting DSCP Values for Egress Packets Based on Flow

Match-any Layer 3 flows may have several match criteria. All flows that match at least one of the match criteria are mapped to the same queue because they are in the same class map.

Setting a DSCP value from QOS-POLICY-IN mode (refer to [Setting a DSCP Value for Egress Packets](#)) assigns the same DSCP value to all of the matching flows in the class-map.

The flow-based DSCP marking feature allows you to assign *different* DSCP to each match criteria.

- Create matching flows within a class map that have different DSCP values.
- CLASS MAP mode

```
match set-ip-dscp
```

The values you set from CLASS-MAP mode override the value you set in the QoS input policy DSCP value. Packets matching the rule are marked with the specified value.

Example of Marking Flows in the Same Queue with Different DSCP Values

```
Dell#show run class-map
!
class-map match-any example-flowbased-dscp
  match ip access-group test set-ip-dscp 2
  match ip access-group test1 set-ip-dscp 4
  match ip precedence 7 set-ip-dscp 1
```

```
Dell#show run qos-policy-input
!
qos-policy-input flowbased
  set ip-dscp 3
```

Displaying Configured Class Maps and Match Criteria

To display all class-maps or a specific class map, use the following command.

Dell Networking OS Behavior: An explicit "deny any" rule in a Layer 3 ACL used in a (match any or match all) class-map creates a "default to Queue 0" entry in the CAM, which causes unintended traffic classification. In the following example, traffic is classified in two Queues, 1 and 2. Class-map ClassAF1 is "match any," and ClassAF2 is "match all".

- Display all class-maps or a specific class map.
- EXEC Privilege mode

```
show qos class-map
```

Examples of Traffic Classifications

The following example shows incorrect traffic classifications.

```
Dell#show running-config policy-map-input
!
policy-map-input PolicyMapIn
  service-queue 1 class-map ClassAF1 qos-policy QosPolicyIn-1
  service-queue 2 class-map ClassAF2 qos-policy QosPolicyIn-2
Dell#show running-config class-map
!
class-map match-any ClassAF1
  match ip access-group AF1-FB1 set-ip-dscp 10
  match ip access-group AF1-FB2 set-ip-dscp 12
  match ip dscp 10 set-ip-dscp 14
  match ipv6 dscp 20 set-ip-dscp 14
!
class-map match-all ClassAF2
  match ip access-group AF2
  match ip dscp 18

Dell#show running-config ACL
!
ip access-list extended AF1-FB1
  seq 5 permit ip host 23.64.0.2 any
  seq 10 deny ip any any
!
ip access-list extended AF1-FB2
  seq 5 permit ip host 23.64.0.3 any
  seq 10 deny ip any any
!
ip access-list extended AF2
  seq 5 permit ip host 23.64.0.5 any
  seq 10 deny ip any any

Dell#show cam layer3-qos interface tengigabitethernet 2/49
Cam  Port Dscp Proto Tcp  Src  Dst SrcIp          DstIp          DSCP   Queue
Index                               Flag Port Port          DstIp          Marking
-----
20416 1    18    IP    0x0  0    0    23.64.0.5/32    0.0.0.0/0     20    2
20417 1    18    IP    0x0  0    0    0.0.0.0/0       0.0.0.0/0     -    0
20418 1    0     IP    0x0  0    0    23.64.0.2/32    0.0.0.0/0     10    1
20419 1    0     IP    0x0  0    0    0.0.0.0/0       0.0.0.0/0     -    0
20420 1    0     IP    0x0  0    0    23.64.0.3/32    0.0.0.0/0     12    1
20421 1    0     IP    0x0  0    0    0.0.0.0/0       0.0.0.0/0     -    0
20422 1    10    0     0x0  0    0    0.0.0.0/0       0.0.0.0/0     14    1
24511 1    0     0     0x0  0    0    0.0.0.0/0       0.0.0.0/0     -    0
```

In the previous example, the ClassAF1 does not classify traffic as intended. Traffic matching the first match criteria is classified to Queue 1, but all other traffic is classified to Queue 0 as a result of CAM entry 20419.

When you remove the explicit "deny any" rule from all three ACLs, the CAM reflects exactly the desired classification.

The following example shows correct traffic classifications.

```
Dell#show cam layer3-qos interface tengigabitethernet 2/49
Cam  Port Dscp Proto Tcp  Src  Dst SrcIp          DstIp          DSCP   Queue
Index                               Flag Port Port          DstIp          Marking
-----
20416 1    18    IP    0x0  0    0    23.64.0.5/32    0.0.0.0/0     20    2
20417 1    0     IP    0x0  0    0    23.64.0.2/32    0.0.0.0/0     10    1
```



20418	1	0	IP	0x0	0	0	23.64.0.3/32	0.0.0.0/0	12	1
20419	1	10	0	0x0	0	0	0.0.0.0/0	0.0.0.0/0	14	1
24511	1	0	0	0x0	0	0	0.0.0.0/0	0.0.0.0/0	-	0

Create a QoS Policy

There are two types of QoS policies — input and output.

Input QoS policies regulate Layer 3 and Layer 2 ingress traffic. The regulation mechanisms for input QoS policies are rate policing and setting priority values. There are two types of input QoS policies: Layer 3 and Layer 2.

Output QoS policies regulate egress traffic. The regulation mechanisms for output QoS policies are bandwidth percentage, scheduler strict, rate shaping, and WRED.

-  **NOTE:** When changing a "service-queue" configuration in a QoS policy map, all QoS rules are deleted and re-added automatically to ensure that the order of the rule is maintained. As a result, the Matched Packets value shown in the `show qos statistics` command is reset.
-  **NOTE:** To avoid issues misconfiguration causes, Dell Networking recommends configuring either DCBX or Egress QoS features, but not both simultaneously. If you enable both DCBX and Egress QoS at the same time, the DCBX configuration is applied and unexpected behavior occurs on the Egress QoS.

Creating an Input QoS Policy

To create an input QoS policy, use the following steps.

1. Create a Layer 3 input QoS policy.

CONFIGURATION mode

```
qos-policy-input
```

Create a Layer 2 input QoS policy by specifying the keyword `layer2` after the `qos-policy-input` command.

2. After you create an input QoS policy, do one or more of the following:

[Configuring Policy-Based Rate Policing](#)

[Setting a DSCP Value for Egress Packets](#)

[Setting a dot1p Value for Egress Packets](#)

Configuring Policy-Based Rate Policing

To configure policy-based rate policing, use the following command.

- Configure rate police ingress traffic.

QOS-POLICY-IN mode

```
rate-police
```

Setting a DSCP Value for Egress Packets

You can set the DSCP value for egress packets based on ingress QoS classification. The 6 bits that are used for DSCP are also used to identify the queue in which traffic is buffered.

Example of Setting a DSCP Value for Egress Packets

```
Dell#config
Dell(conf)#qos-policy-input my-input-qos-policy
Dell(conf-qos-policy-in)#set ip-dscp 34

Dell(conf-qos-policy-in)#show config
!
qos-policy-input my-input-qos-policy
  set ip-dscp 34
Dell(conf-qos-policy-in)#end
Dell#
```

Setting a dot1p Value for Egress Packets

To set a dot1p value for egress packets, use the following command.

- Set a dot1p value for egress packets.
QOS-POLICY-IN mode

```
set mac-dot1p
```

Creating an Output QoS Policy

To create an output QoS policy, use the following commands.

1. Create an output QoS policy.
CONFIGURATION mode

```
qos-policy-output
```
2. After you configure an output QoS policy, do one or more of the following:
 - [Configuring Policy-Based Rate Shaping](#)
 - [Allocating Bandwidth to Queue](#)
 - [Configure a Scheduler to Queue](#)
 - [Specifying WRED Drop Precedence](#)

Configuring Policy-Based Rate Shaping

To configure policy-based rate shaping, use the following command.

- Configure rate shape egress traffic.
QOS-POLICY-OUT mode

```
rate-shape
```

Allocating Bandwidth to Queue

The Dell Networking recommends pre-calculating your bandwidth requirements before creating them. Make sure you apply the QoS policy to all the four queues and that the sum of the bandwidths allocated through them is exactly 100.

When you apply the QoS policies through output policy map and if the sum of the bandwidth percentages configured is below or above 100, the actual bandwidth is allocated proportionally. If the sum of allocated bandwidth is less than 100, the unused bandwidth is allotted to un-allocated queues. If the sums of allocated bandwidth exceed 100, then 1% of the bandwidth is derived for unassigned queues from assigned queues.


- Allocate bandwidth to queues.

QOS-POLICY-OUT mode

```
bandwidth-percentage
```

Configure a Scheduler to Queue

By default, the MXL switch schedules packets for egress based on weighted round robin (WRR).

 **NOTE:** The bandwidth and scheduler cannot be configured at the same time. Policy-level scheduler assigned to queue is applied to both unicast and multicast traffic.

Setting DSCP Values for Egress Packets Based on Flow

Match-any Layer 3 flows may have several match criteria. All flows that match at least one of the match criteria are mapped to the same queue because they are in the same class map.

Setting a DSCP value from QOS-POLICY-IN mode (refer to [Setting a DSCP Value for Egress Packets](#)) assigns the same DSCP value to all of the matching flows in the class-map.

The flow-based DSCP marking feature allows you to assign *different* DSCP to each match criteria.

- Create matching flows within a class map that have different DSCP values.

CLASS MAP mode

```
match ip dscp 0-63 set-ip-dscp 0-63
```

The values you set from CLASS-MAP mode override the value you set in the QoS input policy DSCP value. Packets matching the rule are marked with the specified value.

Example of Marking Flows in the Same Queue with Different DSCP Values

```
Dell#show run class-map
!
class-map match-any example-flowbased-dscp
  match ip access-group test set-ip-dscp 2
  match ip access-group test1 set-ip-dscp 4
  match ip precedence 7 set-ip-dscp 1
```

```
Dell#show run qos-policy-input
!
qos-policy-input flowbased
  set ip-dscp 3
```

Dell#

Specifying WRED Drop Precedence

- Specify a WRED profile to yellow and/or green traffic.
QOS-POLICY-OUT mode

```
wred
```

For more information, refer to [Applying a WRED Profile to Traffic](#).

Create Policy Maps

There are two types of policy maps: input and output.

Creating Input Policy Maps

There are two types of input policy-maps: Layer 3 and Layer 2.

1. Create a Layer 3 input policy map.
CONFIGURATION mode

```
policy-map-input
```

Create a Layer 2 input policy map by specifying the keyword `layer2` with the `policy-map-input` command.

2. After you create an input policy map, do one or more of the following:

[Applying a Class-Map or Input QoS Policy to a Queue](#)

[Applying an Input QoS Policy to an Input Policy Map](#)

[Honoring DSCP Values on Ingress Packets](#)

[Honoring dot1p Values on Ingress Packets](#)

[Enabling Fall Back to Trust Diffserve or dot1p](#)

3. Apply the input policy map to an interface.

Applying a Class-Map or Input QoS Policy to a Queue

To apply a class-map or input QoS policy to a queue, use the following command.

- Assign an input QoS policy to a queue.
POLICY-MAP-IN mode

```
service-queue
```

Applying an Input QoS Policy to an Input Policy Map

To apply an input QoS policy to an input policy map, use the following command.

- Apply an input QoS policy to an input policy map.

POLICY-MAP-IN mode

```
policy-aggregate
```

Honoring DSCP Values on Ingress Packets

The Dell Networking OS provides the ability to honor DSCP values on ingress packets using Trust DSCP feature.

The following table lists the standard DSCP definitions and indicates to which queues the Dell Networking OS maps DSCP values. When you configure trust DSCP, the matched packets and matched bytes counters are not incremented in the `show qos statistics` command.

 **NOTE:** Packets with DSCP value of 63 are dropped.

Table 50. Default DSCP to Queue Mapping

DSCP/CP hex range (XXX)xxx	DSCP Definition	Traditional IP Precedence	Internal Queue ID	DSCP/CP decimal
111XXX		Network Control	3	48–63
110XXX		Internetwork Control	3	48–63
101XXX	EF (Expedited Forwarding)	CRITIC/ECP	2	32–47
100XXX	AF4 (Assured Forwarding)	Flash Override	2	32–47
011XXX	AF3	Flash	1	16–31
010XXX	AF2	Immediate	1	16–31
001XXX	AF1	Priority	0	0–15
000XXX	BE (Best Effort)	Best Effort	0	0–15

- Enable the trust DSCP feature.

POLICY-MAP-IN mode

```
trust dffserv
```

Honoring dot1p Values on Ingress Packets

The Dell Networking OS honors dot1p values on ingress packets with the Trust dot1p feature. The following table specifies the queue to which the classified traffic is sent based on the dot1p value.

Table 51. Default dot1p to Queue Mapping

dot1p	Queue ID
0	0
1	0
2	0
3	1
4	2
5	3
6	3
7	3

The dot1p value is also honored for frames on the default VLAN. For more information, refer to [Priority-Tagged Frames on the Default VLAN](#).

- Enable the trust dot1p feature.
POLICY-MAP-IN mode

```
trust dot1p
```

Enabling Fall Back to Trust Diffserve or dot1p

When using QoS service policies with multiple class maps, you can configure the Dell Networking OS to use the incoming DSCP or dot1p marking as a secondary option for packet queuing if no match occurs in the class maps.

When you use class-maps, traffic is matched against each class-map sequentially from first to last. The sequence is based on the priority of the rules, as follows:

1. Rules with lowest priority, or in the absence of a priority configuration.
2. Rules of the next numerically higher queue.

By default, if no match occurs, the packet is queued to the default queue, Queue 0.

In the following configuration, packets are classified to queues using three class maps:

Example of Viewing Packet Classes Based on DSCP Value

```
!  
policy-map-input input-policy  
  service-queue 1 class-map qos-BE1  
  service-queue 3 class-map qos-AF3  
  service-queue 4 class-map qos-AF4  
  trust diffserv fallback  
!  
class-map match-any qos-AF3  
  match ip dscp 24  
  match ip access-group qos-AF3-ACL  
!  
class-map match-any qos-AF4  
  match ip dscp 32  
  match ip access-group qos-AF4-ACL  
!  
class-map match-all qos-BE1
```

```
match ip dscp 0
match ip access-group qos-BE1-ACL
```

The packet classification logic for the configuration shown is as follows:

1. Match packets against `match-any qos-AF4`. If a match exists, queue the packet as AF4 in Queue 4, and if no match exists, go to the next class map.
2. Match packets against `match-any qos-AF3`. If a match exists, queue the packet as AF3 in Queue 3, and if no match exists, go to the next class map.
3. Match packets against `match-all qos-BE1`. If a match exists, queue the packet as BE1, and if no match exists, queue the packets to the default queue, Queue 0.
4. You can optionally classify packets using their DSCP marking, instead of placing packets in Queue 0, if no match occurs. In the above configuration, if no match occurs against `match-all qos-BE1`, the classification logic continues.
5. Queue the packet according to the DSCP marking. The DSCP to Queue mapping is as noted in [Honoring dot1p Values on Ingress Packets](#).

The behavior is similar for `trust dot1p fallback` in a Layer2 input policy map; the dot1p-to-queue mapping is noted in [Honoring dot1p Values on Ingress Packets](#).

To enable fall back to trust diffserve or dot1p, use the following command.

- Classify packets according to their DSCP value as a secondary option in case no match occurs against the configured class maps.
POLICY-MAP-IN mode

```
trust {diffserve | dot1p} fallback
```

Mapping dot1p Values to Service Queues

All traffic is by default mapped to the same queue, Queue 0.

If you honor dot1p on ingress, you can create service classes based the queueing strategy in [Honoring dot1p Values on Ingress Packets](#). You may apply this queueing strategy globally by entering the following command from CONFIGURATION mode.

- All dot1p traffic is mapped to Queue 0 unless you enable `service-class dynamic dot1p` on an interface or globally.
- Layer 2 or Layer 3 service policies supersede dot1p service classes.
- Create service classes.

INTERFACE mode

```
service-class dynamic dot1p
```

Guaranteeing Bandwidth to dot1p-Based Service Queues

To guarantee bandwidth to dot1p-based service queues, use the following command.

Apply this command in the same way as the `bandwidth-percentage` command in an output QoS policy (refer to [Allocating Bandwidth to Queue](#)). The `bandwidth-percentage` command in QOS-POLICY-OUT mode supersedes the `service-class bandwidth-percentage` command.

- Guarantee a minimum bandwidth to queues globally.

CONFIGURATION mode

```
service-class bandwidth-percentage
```

Applying an Input Policy Map to an Interface

To apply an input policy map to an interface, use the following command.

You can apply the same policy map to multiple interfaces, and you can modify a policy map after you apply it.

- You cannot apply an input Layer 2 QoS policy on an interface you also configure with the `vlan-stack access` command.
- If you apply a service policy that contains an ACL to more than one interface, the system uses ACL optimization to conserve CAM space. The ACL optimization behavior detects when an ACL exists in the CAM rather than writing it to the CAM multiple times.

- Apply an input policy map to an interface.

INTERFACE mode

```
service-policy input
```

Specify the keyword `layer2` if the policy map you are applying a Layer 2 policy map; in this case, INTERFACE mode must be in Switchport mode.

Creating Output Policy Maps

1. Create an output policy map.

CONFIGURATION mode

```
policy-map-output
```

2. After you create an output policy map, do one or more of the following:

[Applying an Output QoS Policy to a Queue](#)

[Specifying an Aggregate QoS Policy](#)

[Applying an Output Policy Map to an Interface](#)

3. Apply the policy map to an interface.

Applying an Output QoS Policy to a Queue

To apply an output QoS policy to a queue, use the following command.

- Apply an output QoS policy to queues.

INTERFACE mode

```
service-queue
```

Specifying an Aggregate QoS Policy

To specify an aggregate QoS policy, use the following command.

- Specify an aggregate QoS policy.

POLICY-MAP-OUT mode

```
policy-aggregate
```

Applying an Output Policy Map to an Interface

To apply an output policy map to an interface, use the following command.

- Apply an input policy map to an interface.

INTERFACE mode

```
service-policy output
```

You can apply the same policy map to multiple interfaces, and you can modify a policy map after you apply it.

Enabling QoS Rate Adjustment

By default, while rate limiting, policing, and shaping, the Dell Networking OS does not include the Preamble, SFD, or the IFG fields. These fields are overhead; only the fields from MAC destination address to the CRC are used for forwarding and are included in these rate metering calculations.

The Ethernet packet format consists of:

- Preamble: 7 bytes Preamble
- Start frame delimiter (SFD): 1 byte
- Destination MAC address: 6 bytes
- Source MAC address: 6 bytes
- Ethernet Type/Length: 2 bytes
- Payload: (variable)
- Cyclic redundancy check (CRC): 4 bytes
- Inter-frame gap (IFG): (variable)

You can optionally include overhead fields in rate metering calculations by enabling QoS rate adjustment.

QoS rate adjustment is disabled by default, and `no qos-rate-adjust` is listed in the running-configuration

- Include a specified number of bytes of packet overhead to include in rate limiting, policing, and shaping calculations.

CONFIGURATION mode

```
qos-rate-adjust overhead-bytes
```

For example, to include the Preamble and SFD, enter `qos-rate-adjust 8`. For variable length overhead fields, know the number of bytes you want to include.

The default is disabled.

Enabling Strict-Priority Queueing

Strict-priority means that the Dell Networking OS de-queues all packets from the assigned queue before servicing any other queues.

- The `strict-priority` supersedes `bandwidth-percentage` and `bandwidth-weight percentage` configurations.
- A queue with strict priority can starve other queues in the same port-pipe.
- If more than two strict priority queues are configured, the strict priority queue with a higher queue number is scheduled first.
- Assign strict priority to one unicast queue.

CONFIGURATION mode

```
strict-priority
```

The range is from 1 to 3.

Weighted Random Early Detection

The WRED congestion avoidance mechanism drops packets to prevent buffering resources from being consumed.

Traffic is a mixture of various kinds of packets. The rate at which some types of packets arrive might be greater than others. In this case, the space on the buffer and traffic manager (BTM) (ingress or egress) can be consumed by only one or a few types of traffic, leaving no space for other types. You can apply a WRED profile to a policy-map so that specified traffic can be prevented from consuming too much of the BTM resources.

WRED uses a profile to specify minimum and maximum threshold values. The minimum threshold is the allotted buffer space for specified traffic, for example, 1000KB on egress. If the 1000KB is consumed, packets are dropped randomly at an exponential rate until the maximum threshold is reached (as shown in the following illustration); this procedure is the *early detection* part of WRED. If the maximum threshold, for example, 2000KB, is reached, all incoming packets are dropped until the buffer space consumes less than 2000KB of the specified traffic.

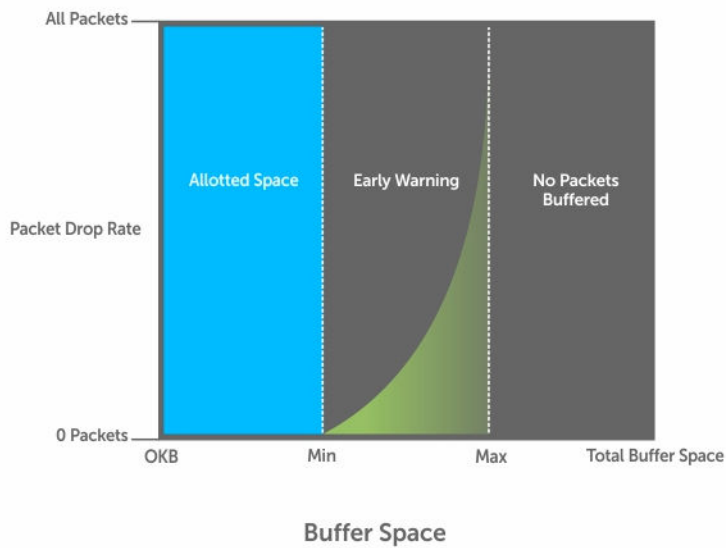


Figure 106. Packet Drop Rate for WRED

You can create a custom WRED profile or use one of the five pre-defined profiles.

Table 52. Pre-Defined WRED Profiles

Default Profile Name	Minimum Threshold	Maximum Threshold	Maximum Drop Rate
wred_drop	0	0	100
wred_teng_y	467	4671	100
wred_teng_g	467	4671	50
wred_fortyg_y	467	4671	50
wred_fortyg_g	467	4671	25

Creating WRED Profiles

To create WRED profiles, use the following commands.

1. Create a WRED profile.
CONFIGURATION mode

wred-profile
2. Specify the minimum and maximum threshold values.
WRED mode

threshold

Applying a WRED Profile to Traffic

After you create a WRED profile, you must specify to which traffic the system should apply the profile. The Dell Networking OS assigns a color (also called drop precedence) — red, yellow, or green — to each packet based on its DSCP value before queuing it.

DSCP is a 6-bit field. Dell Networking uses the first 3 bits of this field (DP) to determine the drop precedence.

- DP values of 110, 100, and 101 map to yellow; all other values map to green.
- If you do not configure the system to honor DSCP values on ingress (refer to [Honoring DSCP Values on Ingress Packets](#)), all traffic defaults to green drop precedence.
- Assign a WRED profile to either yellow or green traffic.
QOS-POLICY-OUT mode

```
wred
```

Displaying Default and Configured WRED Profiles

To display the default and configured WRED profiles, use the following command.

- Display default and configured WRED profiles and their threshold values.
EXEC mode

```
show qos wred-profile
```

Displaying WRED Profiles

```
Dell#show qos wred-profile
```

```
Wred-profile-name min-threshold max-threshold max-drop-rate
wred_drop          0              0              100
wred_teng_y        467            4671           100
wred_teng_g        467            4671           50
0
Dell#
```

Displaying WRED Drop Statistics

To display WRED drop statistics, use the following command.

- Display the number of packets the system the WRED profile drops.
EXEC Privilege mode

```
show qos statistics
```

Example of the `show qos statistics wred-profile` Command

```
Dell#show qos statistics wred-profile
```

```
Interface Te 0/20
```

```
Drop-statistic Dropped Pkts
Green          11234
Yellow         12484
Out of Profile 0
```

```
Dell#
```

Displaying egress-queue Statistics

To display egress-queue statistics of both transmitted and dropped packets and bytes, use the following command.

- Display the number of packets and number of bytes on the egress-queue profile.

EXEC Privilege mode

```
show qos statistics egress-queue
```

Example of the `show qos statistics egress-queue` Command

```
Dell#show qos statistics egress-queue
```

```
Interface Te 1/1
Unicast/Multicast Egress Queue Statistics
Queue# Q# Type TxPkts TxBytes DroppedPkts DroppedBytes
-----
0 UCAST 0 0 0 0
1 UCAST 0 0 0 0
2 UCAST 0 0 0 0
3 UCAST 0 0 0 0
4 UCAST 0 0 0 0
5 UCAST 0 0 0 0
6 UCAST 0 0 0 0
7 UCAST 0 0 0 0
8 UCAST 204 13056 0 0
9 MCAST 0 0 0 0
10 MCAST 0 0 0 0
11 MCAST 0 0 0 0
12 MCAST 0 0 0 0
13 MCAST 0 0 0 0
14 MCAST 0 0 0 0
15 MCAST 0 0 0 0
16 MCAST 0 0 0 0
17 MCAST 0 0 0 0
```

```
Dell#
```

Classifying Layer 2 Traffic on Layer 3 Interfaces

To process Layer 3 packets that contain Dot1p — (IEEE 802.1p) Packet classification (Layer 2 headers), configure VLAN tags on a physical Layer 3 interface (that is configured with an IP address and is not associated with any VLAN). You can also configure a VLAN subinterface over a physical underlying interface and classify packets using the dot1p value.

To apply an input policy map to Layer 3 physical interfaces, use the `service-policy input policy-name layer 2` command in Interface Configuration mode.

To apply a Layer 2 policy on Layer 3 interfaces, perform the following:

1. Configure an interface with an IP address or a VLAN subinterface
CONFIGURATION mode

```
Dell(conf)# int fo 0/0
```

INTERFACE mode

```
Dell(conf-if-fo-0/0)# ip address 90.1.1.1/16
```

2. Configure the Layer 2 policy with Layer 2 (Dot1p or source MAC-based) classification rules.
CONFIGURATION mode

```
Dell(conf)# policy-map-input l2p layer2
```

3. Apply the Layer 2 policy on the Layer 3 interface.

INTERFACE mode

```
Dell(conf-if-fo-0/0)# service-policy input l2p layer2
```

Classifying Packets Based on a Combination of DSCP Code Points and VLAN IDs

You can configure a classifier map, which contains both the Differentiated Services Code Point (DSCP) and MAC VLAN IDs as parameters, for filtering packets that are received before they are forwarded or dropped. You can now specify both DSCP-IP packet classification (Layer 3 headers) and Dot1p—(IEEE 802.1p) Packet classification (Layer 2 headers) as match criteria in a Layer 3 class map.

The type of the class map is determined during the creation of a class map. In releases of Dell Networking OS earlier than Release 9.2(0.0), you can configure only the dot1p value as the filter criterion in Layer 2 class maps and the DSCP value as the filter parameter in Layer 3 class maps. It was also possible to classify packets using both the Layer 2 attribute, dot1p value or MAC VLAN, in a Layer 2 class map and the Layer 3 attribute, DSCP value, in a Layer 3 class map. However, it was not possible to configure both dot1p or MAC VLAN, and DSCP values in the same Layer 2 or Layer 3 class map.

All class maps are Layer 3 by default. You can now configure a Layer 3 class map to differentiate traffic according to the IP VLAN value and the DSCP value. You can use the `match ip vlan vlan-id` command in Class Map Input Configuration mode to specify a match criterion for a class map based on a VLAN ID. You can attach this class map with a policy map, and associate the policy map with a service queue. When you link class maps to queues using the `service-queue` command, Dell Networking OS matches the class-maps according to queue priority (queue numbers closer to 0 have lower priorities).

To create IP VLAN and DSCP values as match criteria in a Layer 3 class map, and to associate the class map with a policy map that is linked to a service queue, perform the following:

1. Create a match-any or a match-all Layer 3 class map, depending on whether you want the packets to meet all or any of the match criteria to be a member of the class. By default, a Layer 3 class map is created if you do not enter the `layer2` option with the `class-map` command. When you create a class map, you enter the Class Map configuration mode.

CONFIGURATION mode

```
Dell (conf)#class-map match-all pp_classmap
```

2. Use a DSCP value as a match criterion.

CLASS-MAP mode

```
Dell (conf-class-map)#match ip dscp 5
```

3. Configure a match criterion for a class map based on the VLAN ID.

CLASS-MAP mode

```
Dell (conf-class-map)#match ip vlan 5
```

4. Create a QoS input policy on the device.

CONFIGURATION mode

```
Dell(conf)#qos-policy-input pp_qospolicy
```

5. Specify the DSCP value to be set on the matched traffic.

QOS-POLICY-IN mode

```
Dell(conf-qos-policy-in)#set ip-dscp 5
```

6. Create an input policy map.

CONFIGURATION mode

```
Dell(conf)#policy-map-input pp_policmap
```

7. Create a service queue to associate the class map and QoS policy map.

POLICY-MAP mode

```
Dell(conf-policy-map-in)#service-queue 0 class-map pp_classmap qos-policy  
pp_qospolicy
```

Routing Information Protocol (RIP)

The routing information protocol (RIP) is based on a distance-vector algorithm and tracks distances or hop counts to nearby routers when establishing network connections.

RIP protocol standards are listed in the [Standards Compliance](#) chapter.

Protocol Overview

RIP is the oldest interior gateway protocol.

There are two versions of RIP: RIP version 1 (RIPv1) and RIP version 2 (RIPv2). These versions are documented in RFCs 1058 and 2453.

RIPv1

RIPv1 learns where nodes in a network are located by automatically constructing a routing data table.

The routing table is established after RIP sends out one or more broadcast signals to all adjacent nodes in a network. Hop counts of these signals are tracked and entered into the routing table, which defines where nodes in the network are located.

The information that is used to update the routing table is sent as either a request or response message. In RIPv1, automatic updates to the routing table are performed as either one-time requests or periodic responses (every 30 seconds). RIP transports its responses or requests by means of user datagram protocol (UDP) over port 520.

RIP must receive regular routing updates to maintain a correct routing table. Response messages containing a router's full routing table are transmitted every 30 seconds. If a router does not send an update within a certain amount of time, the hop count to that route is changed to unreachable (a route hop metric of 16 hops). Another timer sets the amount of time before the unreachable routes are removed from the routing table.

This first RIP version does not support variable length subnet mask (VLSM) or classless inter-domain routing (CIDR) and is not widely used.

RIPv2

RIPv2 adds support for subnet fields in the RIP routing updates, thus qualifying it as a classless routing protocol.

The RIPv2 message format includes entries for route tags, subnet masks, and next hop addresses. Another enhancement included in RIPv2 is multicasting for route updates on IP multicast address 224.0.0.9.

Implementation Information

The Dell Networking OS supports both versions of RIP and allows you to configure one version globally and the other version on interfaces or both versions on the interfaces.

The following table lists the defaults for RIP in the system.

Table 53. RIP Defaults

Feature	Default
Interfaces running RIP	<ul style="list-style-type: none">• Listen to RIPv1 and RIPv2• Transmit RIPv1
RIP timers	<ul style="list-style-type: none">• update timer = 30 seconds• invalid timer = 180 seconds• holddown timer = 180 seconds• flush timer = 240 seconds
Auto summarization	Enabled
ECMP paths supported	16

Configuration Information

By default, RIP is disabled in the system.

To configure RIP, you must use commands in two modes: ROUTER RIP and INTERFACE. Commands executed in the ROUTER RIP mode configure RIP globally, while commands executed in the INTERFACE mode configure RIP features on that interface only.

RIP is best suited for small, homogeneous networks. You must configure all devices within the RIP network to support RIP if they are to participate in the RIP.

Configuration Task List

The following is the configuration task list for RIP.

- [Enabling RIP Globally](#) (mandatory)
- [Configure RIP on Interfaces](#) (optional)

- [Controlling RIP Routing Updates](#) (optional)
- [Setting the Send and Receive Version](#) (optional)
- [Generating a Default Route](#) (optional)
- [Controlling Route Metrics](#) (optional)
- [Summarize Routes](#) (optional)
- [Controlling Route Metrics](#)
- [Debugging RIP](#)

For a complete listing of all commands related to RIP, refer to the *Dell Networking OS Command Reference Interface Guide*.

Enabling RIP Globally

By default, RIP is not enabled in the system.

To enable RIP globally, use the following commands.

1. Enter ROUTER RIP mode and enable the RIP process on the system.

CONFIGURATION mode

```
router rip
```

2. Assign an IP network address as a RIP network to exchange routing information.

ROUTER RIP mode

```
network ip-address
```

Example of Verifying RIP is Enabled

Example of Viewing Learned RIP Routes

After designating networks with which the system is to exchange RIP information, ensure that all devices on that network are configured to exchange RIP information.

The Dell Networking OS default is to send RIPv1 and to receive RIPv1 and RIPv2. To change the RIP version globally, use the `version` command in ROUTER RIP mode.

To view the global RIP configuration, use the `show running-config` command in EXEC mode or the `show config` command in ROUTER RIP mode.

```
Dell(conf-router_rip)#show config
!
router rip
  network 10.0.0.0
Dell(conf-router_rip)#
```

When the RIP process has learned the RIP routes, use the `show ip rip database` command in EXEC mode to view those routes.

```
Dell#show ip rip database
Total number of routes in RIP database: 978
160.160.0.0/16
  [120/1] via 29.10.10.12, 00:00:26, Fa 0/0
160.160.0.0/16    auto-summary
2.0.0.0/8
  [120/1] via 29.10.10.12, 00:01:22, Fa 0/0
```

```

2.0.0.0/8          auto-summary
4.0.0.0/8
  [120/1] via 29.10.10.12, 00:01:22, Fa 0/0
4.0.0.0/8          auto-summary
8.0.0.0/8
  [120/1] via 29.10.10.12, 00:00:26, Fa 0/0
8.0.0.0/8          auto-summary
12.0.0.0/8
  [120/1] via 29.10.10.12, 00:00:26, Fa 0/0
12.0.0.0/8          auto-summary
20.0.0.0/8
  [120/1] via 29.10.10.12, 00:00:26, Fa 0/0
20.0.0.0/8          auto-summary
29.10.10.0/24      directly connected,Fa 0/0
29.0.0.0/8          auto-summary
31.0.0.0/8
  [120/1] via 29.10.10.12, 00:00:26, Fa 0/0
31.0.0.0/8          auto-summary
192.162.2.0/24
  [120/1] via 29.10.10.12, 00:01:21, Fa 0/0
192.162.2.0/24      auto-summary
192.161.1.0/24
  [120/1] via 29.10.10.12, 00:00:27, Fa 0/0
192.161.1.0/24      auto-summary
192.162.3.0/24
  [120/1] via 29.10.10.12, 00:01:22, Fa 0/0
192.162.3.0/24      auto-summary

```

To disable RIP globally, use the `no router rip` command in CONFIGURATION mode.

Configure RIP on Interfaces

When you enable RIP globally on the system, interfaces meeting certain conditions start receiving RIP routes.

By default, interfaces that you enable and configure with an IP address in the same subnet as the RIP network address receive RIPv1 and RIPv2 routes and send RIPv1 routes.

Assign IP addresses to interfaces that are part of the same subnet as the RIP network identified in the `network` command syntax.

Controlling RIP Routing Updates

By default, RIP broadcasts routing information out all enabled interfaces, but you can configure RIP to send or to block RIP routing information, either from a specific IP address or a specific interface. To control which devices or interfaces receive routing updates, configure a direct update to one router and configure interfaces to block RIP updates from other sources.

To control the source of RIP route information, use the following commands.

- Define a specific router to exchange RIP information between it and the Dell Networking system.
ROUTER RIP mode

```
neighbor ip-address
```

You can use this command multiple times to exchange RIP information with as many RIP networks as you want.

- Disable a specific interface from sending or receiving RIP routing information.

ROUTER RIP mode

```
passive-interface interface
```

Adding RIP Routes from Other Instances

In addition to filtering routes, you can add routes from other routing instances or protocols to the RIP process.

With the `redistribute` command, you can include open shortest path first (OSPF), static, or directly connected routes in the RIP process.

To add routes from other routing instances or protocols, use the following commands.

- Include directly connected or user-configured (static) routes in RIP.

ROUTER RIP mode

```
redistribute {connected | static} [metric metric-value] [route-map map-name]
```

- *metric-value*: the range is from 0 to 16.
- *map-name*: the name of a configured route map.

- Include specific OSPF routes in RIP.

ROUTER RIP mode

```
redistribute ospf process-id [match external {1 | 2} | match internal]  
[metric value] [route-map map-name]
```

Configure the following parameters:

- *process-id*: the range is from 1 to 65535.
- *metric*: the range is from 0 to 16.
- *map-name*: the name of a configured route map.

To view the current RIP configuration, use the `show running-config` command in EXEC mode or the `show config` command in ROUTER RIP mode.

Assigning a Prefix List to RIP Routes

Another method of controlling RIP (or any routing protocol) routing information is to filter the information through a prefix list. A prefix list is applied to incoming or outgoing routes.

Those routes must meet the conditions of the prefix list; if not, the system drops the route. Prefix lists are globally applied on all interfaces running RIP. Configure the prefix list in PREFIX LIST mode prior to assigning it to the RIP process.

For configuration information about prefix lists, refer to [Access Control Lists \(ACLs\)](#).

To apply prefix lists to incoming or outgoing RIP routes, use the following commands.

- Assign a configured prefix list to all incoming RIP routes.

ROUTER RIP mode

```
distribute-list prefix-list-name in
```

- Assign a configured prefix list to all outgoing RIP routes.

ROUTER RIP mode

```
distribute-list prefix-list-name out
```

To view the current RIP configuration, use the `show running-config` command in EXEC mode or the `show config` command in ROUTER RIP mode.

Setting the Send and Receive Version

To change the RIP version globally or on an interface in the system, use the following command. To specify the RIP version, use the `version` command in ROUTER RIP mode. To set an interface to receive only one or the other version, use the `ip rip send version` or the `ip rip receive version` commands in INTERFACE mode.

You can set one RIP version globally on the system using `system`. This command sets the RIP version for RIP traffic on the interfaces participating in RIP unless the interface was specifically configured for a specific RIP version.

- Set the RIP version sent and received on the system.

ROUTER RIP mode

```
version {1 | 2}
```

- Set the RIP versions received on that interface.

INTERFACE mode

```
ip rip receive version [1] [2]
```

- Set the RIP versions sent out on that interface.

INTERFACE mode

```
ip rip send version [1] [2]
```

Example of an Interface Participating in the RIP Process

Example of Configuring an Interface to Send/Receive Specified Versions of RIP

Example of the `show ip protocols` Command to Verify RIP Versions on an Interface

To see whether the `version` command is configured, use the `show config` command in ROUTER RIP mode. To view the routing protocols configuration, use the `show ip protocols` command in EXEC mode.

The following example shows the RIP configuration after the ROUTER RIP mode `version` command is set to RIPv2. When you set the ROUTER RIP mode `version` command, the interface (GigabitEthernet 0/0) participating in the RIP process is also set to send and receive RIPv2 (shown in bold).

```
Dell#show ip protocols
```

```
Routing Protocols is RIP
Sending updates every 30 seconds, next due in 23
Invalid after 180 seconds, hold down 180, flushed after 240
Output delay 8 milliseconds between packets
Automatic network summarization is in effect
Outgoing filter for all interfaces is
Incoming filter for all interfaces is
Default redistribution metric is 1
Default version control: receive version 2, send version 2
  Interface      Recv  Send
  GigabitEthernet 0/0  2    2
Routing for Networks:
  10.0.0.0

Routing Information Sources:
```



```
Gateway      Distance      Last Update
Distance: (default is 120)
```

Dell#

To configure an interface to receive or send both versions of RIP, include 1 and 2 in the command syntax. The command syntax for sending both RIPv1 and RIPv2 and receiving only RIPv2 is shown in the following example.

```
Dell(conf-if)#ip rip send version 1 2
Dell(conf-if)#ip rip receive version 2
```

The following example of the `show ip protocols` command confirms that both versions are sent out that interface. This interface no longer sends and receives the same RIP versions as the Dell Networking OS does globally (shown in bold).

Dell#show ip protocols

```
Routing Protocols is RIP
Sending updates every 30 seconds, next due in 11
Invalid after 180 seconds, hold down 180, flushed after 240
Output delay 8 milliseconds between packets
Automatic network summarization is in effect
Outgoing filter for all interfaces is
Incoming filter for all interfaces is
Default redistribution metric is 1
Default version control: receive version 2, send version 2
  Interface      Recv  Send
  FastEthernet 0/0  2    1 2
Routing for Networks:
  10.0.0.0
```

```
Routing Information Sources:
  Gateway      Distance      Last Update
Distance: (default is 120)
```

Dell#

Generating a Default Route

Traffic is forwarded to the default route when the traffic's network is not explicitly listed in the routing table.

Default routes are not enabled in RIP unless specified. Use the `default-information originate` command in ROUTER RIP mode to generate a default route into RIP. In the Dell Networking OS, default routes received in RIP updates from other routes are advertised if you configure the `default-information originate` command.

- Specify the generation of a default route in RIP.
ROUTER RIP mode

```
default-information originate [always] [metric value] [route-map route-map-name]
```

- `always`: Enter the keyword `always` to always generate a default route.

- *value* The range is from 1 to 16.
- *route-map-name*: The name of a configured route map.

To confirm that the default route configuration is completed, use the `show config` command in ROUTER RIP mode.


Summarize Routes

Routes in the RIPv2 routing table are summarized by default, thus reducing the size of the routing table and improving routing efficiency in large networks.

By default, the `autosummary` command in ROUTER RIP mode is enabled and summarizes RIP routes up to the classful network boundary.

If you must perform routing between discontinuous subnets, disable automatic summarization. With automatic route summarization disabled, subnets are advertised.

The `autosummary` command requires no other configuration commands. To disable automatic route summarization, enter `no autosummary` in ROUTER RIP mode.

 **NOTE:** If you enable the `ip split-horizon` command on an interface, the system does not advertise the summarized address.

Controlling Route Metrics

As a distance-vector protocol, RIP uses hop counts to determine the best route, but sometimes the shortest hop count is a route over the lowest-speed link.

To manipulate RIP routes so that the routing protocol prefers a different route, manipulate the route by using the `offset` command.

Exercise caution when applying an `offset` command to routers on a broadcast network, as the router using the `offset` command is modifying RIP advertisements before sending out those advertisements.

The `distance` command also allows you to manipulate route metrics. To assign different weights to routes so that the ones with the lower weight or administrative distance assigned are preferred, use the `distance` command.

To set route matrixes, use the following commands.

- Apply a weight to all routes or a specific route and ACL.

ROUTER RIP mode

```
distance weight [ip-address mask [access-list-name]]
```

Configure the following parameters:

- *weight*: the range is from 1 to 255. The default is **120**.
- *ip-address mask*: the IP address in dotted decimal format (A.B.C.D), and the mask in slash format (/x).
- *access-list-name*: the name of a configured IP ACL.
- Apply an additional number to the incoming or outgoing route metrics.

ROUTER RIP mode

```
offset-list access-list-name {in | out} offset [interface]
```

Configure the following parameters:

- *prefix-list-name*: the name of an established Prefix list to determine which incoming routes are modified
- *offset*: the range is from 0 to 16.
- *interface*: the type, slot, and number of an interface.

To view the configuration changes, use the `show config` command in ROUTER RIP mode.

Debugging RIP

The `debug ip rip` command enables RIP debugging.

When you enable debugging, you can view information on RIP protocol changes or RIP routes.

To enable RIP debugging, use the following command.

- `debug ip rip [interface | database | events | trigger]`
EXEC privilege mode

Enable debugging of RIP.

Example of the `debug ip rip` Command

The following example shows the confirmation when you enable the debug function.

```
Dell#debug ip rip
RIP protocol debug is ON
Dell#
```

To disable RIP, use the `no debug ip rip` command.

RIP Configuration Example

The examples in this section show the command sequence to configure RIPv2 on the two routers shown in the following illustration — *Core 2* and *Core 3*.

The host prompts used in the following example reflect those names. The examples are divided into the following groups of command sequences:

- [RIP Configuration on Core2](#)
- [Core 2 RIP Output](#)
- [RIP Configuration on Core3](#)
- [Core 3 RIP Output](#)
- [RIP Configuration Summary](#)




Figure 107. RIP Topology Example

RIP Configuration on Core2

The following example shows how to configure RIPv2 on a host named Core2.

Example of Configuring RIPv2 on Core 2

```
Core2(conf-if-gi-2/31)#
Core2(conf-if-gi-2/31)#router rip
Core2(conf-router_rip)#ver 2
Core2(conf-router_rip)#network 10.200.10.0
Core2(conf-router_rip)#network 10.300.10.0
Core2(conf-router_rip)#network 10.11.10.0
Core2(conf-router_rip)#network 10.11.20.0
Core2(conf-router_rip)#show config
!
router rip
  network 10.0.0.0
  version 2
Core2(conf-router_rip)#
```

Core 2 RIP Output

The examples in the section show the core 2 RIP output.

Example of the show ip rip database Command to View Learned RIP Routes on Core 2

Example of the show ip route Command to Show RIP Setup on Core 2

Example of the show ip protocols Command to Show RIP Configuration Activity on Core 2

- To display Core 2 RIP database, use the show ip rip database command.
- To display Core 2 RIP setup, use the show ip route command.
- To display Core 2 RIP activity, use the show ip protocols command.

```
Core2(conf-router_rip)#end
00:12:24: %RPM0-P:CP %SYS-5-CONFIG_I: Configured from console by console
Core2#show ip rip database
Total number of routes in RIP database: 7
10.11.30.0/24
  [120/1] via 10.11.20.1, 00:00:03, TenGigabitEthernet 2/31
10.300.10.0/24   directly connected,TenGigabitEthernet 2/42
10.200.10.0/24   directly connected,TenGigabitEthernet 2/41
10.11.20.0/24    directly connected,TenGigabitEthernet 2/31
10.11.10.0/24    directly connected,TenGigabitEthernet 2/11
10.0.0.0/8       auto-summary
192.168.1.0/24
```

```

    [120/1] via 10.11.20.1, 00:00:03, TenGigabitEthernet 2/31
192.168.1.0/24      auto-summary
192.168.2.0/24
    [120/1] via 10.11.20.1, 00:00:03, TenGigabitEthernet 2/31
192.168.2.0/24      auto-summary
Core2#

```

Core2#show ip route

```

Codes: C - connected, S - static, R - RIP,
       B - BGP, IN - internal BGP, EX - external BGP, LO - Locally Originated,
       O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
       E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,
       L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,
       > - non-active route, + - summary route

```

Gateway of last resort is not set

Destination	Gateway	Dist/Metric	Last Change
C	10.11.10.0/24	Direct, TenGig 2/11	0/0 00:02:26
C	10.11.20.0/24	Direct, TenGig 2/31	0/0 00:02:02
R	10.11.30.0/24	via 10.11.20.1, TenGig 2/31	120/1 00:01:20
C	10.200.10.0/24	Direct, TenGig 2/41	0/0 00:03:03
C	10.300.10.0/24	Direct, TenGig 2/42	0/0 00:02:42
R	192.168.1.0/24	via 10.11.20.1, TenGig 2/31	120/1 00:01:20
R	192.168.2.0/24	via 10.11.20.1, TenGig 2/31	120/1 00:01:20
Core2#			
R	192.168.1.0/24	via 10.11.20.1, TenGig 2/31	120/1 00:05:22
R	192.168.2.0/24	via 10.11.20.1, TenGig 2/31	120/1 00:05:22

Core2#

Core2#show ip protocols

```

Routing Protocol is "RIP"
  Sending updates every 30 seconds, next due in 17
  Invalid after 180 seconds, hold down 180, flushed after 240
  Output delay 8 milliseconds between packets
  Automatic network summarization is in effect
  Outgoing filter for all interfaces is
  Incoming filter for all interfaces is
  Default redistribution metric is 1
  Default version control: receive version 2, send version 2
  Interface Recv Send
  TenGigabitEthernet 2/42 2 2
  TenGigabitEthernet 2/41 2 2
  TenGigabitEthernet 2/31 2 2
  TenGigabitEthernet 2/11 2 2

```

Routing for Networks:

```

10.300.10.0
10.200.10.0
10.11.20.0
10.11.10.0

```

Routing Information Sources:

Gateway	Distance	Last Update
10.11.20.1	120	00:00:12

Distance: (default is 120)

Core2#

RIP Configuration on Core3

The following example shows how to configure RIPv2 on a host named Core3.

Example of Configuring RIPv2 on Core3

```
Core3(conf-if-gi-3/21)#router rip
Core3(conf-router_rip)#version 2
Core3(conf-router_rip)#network 192.168.1.0
Core3(conf-router_rip)#network 192.168.2.0
Core3(conf-router_rip)#network 10.11.30.0
Core3(conf-router_rip)#network 10.11.20.0
Core3(conf-router_rip)#show config
!
router rip
 network 10.0.0.0
 network 192.168.1.0
 network 192.168.2.0
 version 2
Core3(conf-router_rip)#
```

Core 3 RIP Output

The examples in this section show the core 2 RIP output.

- To display Core 3 RIP database, use the `show ip rip database` command.
- To display Core 3 RIP setup, use the `show ip route` command.
- To display Core 3 RIP activity, use the `show ip protocols` command.

Example of the `show ip rip database` Command to View Learned RIP Routes on Core 3

Example of the `show ip routes` Command to View RIP Setup on Core 3

Example of the `show ip protocols` Command to Show RIP Configuration Activity on Core 3

```
Core3#show ip rip database
Total number of routes in RIP database: 7
10.11.10.0/24
 [120/1] via 10.11.20.2, 00:00:13, TenGigabitEthernet 3/21
10.200.10.0/24
 [120/1] via 10.11.20.2, 00:00:13, TenGigabitEthernet 3/21
10.300.10.0/24
 [120/1] via 10.11.20.2, 00:00:13, TenGigabitEthernet 3/21
10.11.20.0/24    directly connected,TenGigabitEthernet 3/21
10.11.30.0/24   directly connected,TenGigabitEthernet 3/11
10.0.0.0/8      auto-summary
192.168.1.0/24  directly connected,TenGigabitEthernet 3/43
192.168.1.0/24  auto-summary
192.168.2.0/24  directly connected,TenGigabitEthernet 3/44
192.168.2.0/24  auto-summary
Core3#
```

```
Core3#show ip routes
```

```
Codes: C - connected, S - static, R - RIP,
        B - BGP, IN - internal BGP, EX - external BGP, LO - Locally Originated,
        O - OSPF, IA - OSPF inter area, N1 - OSPF NSSA external type 1,
        N2 - OSPF NSSA external type 2, E1 - OSPF external type 1,
        E2 - OSPF external type 2, i - IS-IS, L1 - IS-IS level-1,
```

L2 - IS-IS level-2, IA - IS-IS inter area, * - candidate default,
> - non-active route, + - summary route

```
Gateway of last resort is not set
  Destination      Gateway             Dist/Metric  Last Change
  -----
R 10.11.10.0/24    via 10.11.20.2, TenGig 3/21    120/1      00:01:14
C 10.11.20.0/24    Direct, TenGig 3/21    0/0        00:01:53
C 10.11.30.0/24    Direct, TenGig 3/11    0/0        00:06:00
R 10.200.10.0/24   via 10.11.20.2, TenGig 3/21 120/1      00:01:14
R 10.300.10.0/24   via 10.11.20.2, TenGig 3/21 120/1      00:01:14
C 192.168.1.0/24   Direct, TenGig 3/43 0/0        00:06:53
C 192.168.2.0/24   Direct, TenGig 3/44 0/0        00:06:26
Core3#
```

```
Core3#show ip protocols
```

```
Routing Protocol is "RIP"
  Sending updates every 30 seconds, next due in 6
  Invalid after 180 seconds, hold down 180, flushed after 240
  Output delay 8 milliseconds between packets
  Automatic network summarization is in effect
  Outgoing filter for all interfaces is
  Incoming filter for all interfaces is
  Default redistribution metric is 1
  Default version control: receive version 2, send version 2
    Interface Recv Send
    TenGigabitEthernet 3/21 2 2
    TenGigabitEthernet 3/11 2 2
    TenGigabitEthernet 3/44 2 2
    TenGigabitEthernet 3/43 2 2
Routing for Networks:
  10.11.20.0
  10.11.30.0
  192.168.2.0
  192.168.1.0
```

```
Routing Information Sources:
  Gateway      Distance  Last Update
  10.11.20.2   120      00:00:22
```

```
Distance: (default is 120)
```

```
Core3#
```

RIP Configuration Summary

Example of Viewing RIP Configuration on Core 2

Example of Viewing RIP Configuration on Core 3

```
!
interface TenGigabitEthernet 2/11
  ip address 10.11.10.1/24
  no shutdown
!
interface TenGigabitEthernet 2/31
  ip address 10.11.20.2/24
  no shutdown
!
interface TenGigabitEthernet 2/41
  ip address 10.200.10.1/24
  no shutdown
```

```
!  
interface TenGigabitEthernet 2/42  
  ip address 10.250.10.1/24  
  no shutdown  
  
router rip  
version 2  
10.200.10.0  
10.300.10.0  
10.11.10.0  
10.11.20.0  
  
!  
interface TenGigabitEthernet 3/11  
  ip address 10.11.30.1/24  
  no shutdown  
  
!  
interface TenGigabitEthernet 3/21  
  ip address 10.11.20.1/24  
  no shutdown  
  
!  
interface TenGigabitEthernet 3/43  
  ip address 192.168.1.1/24  
  no shutdown  
  
!  
interface TenGigabitEthernet 3/44  
  ip address 192.168.2.1/24  
  no shutdown  
  
!  
router rip  
version 2  
network 10.11.20.0  
network 10.11.30.0  
network 192.168.1.0  
network 192.168.2.0
```


Remote Monitoring (RMON)

RMON is an industry-standard implementation that monitors network traffic by sharing network monitoring information.

RMON provides both 32-bit and 64-bit monitoring facility and long-term statistics collection on Dell Networking Ethernet interfaces.

RMON operates with the simple network management protocol (SNMP) and monitors all nodes on a local area network (LAN) segment. RMON monitors traffic passing through the router and segment traffic not destined for the router. The monitored interfaces may be chosen by using alarms and events with standard management information bases (MIBs).

Implementation Information

Configure SNMP prior to setting up RMON.

For a complete SNMP implementation description, refer to [Simple Network Management Protocol \(SNMP\)](#).

Configuring RMON requires using the RMON CLI and includes the following tasks:

- [Setting the rmon Alarm](#)
- [Configuring an RMON Event](#)
- [Configuring RMON Collection Statistics](#)
- [Configuring the RMON Collection History](#)
- [Enabling an RMON MIB Collection History Group](#)


RMON implements the following standard request for comments (RFCs) (for more information, refer to the [Standards Compliance](#) chapter).

- RFC-2819
- RFC-3273
- RFC-3434
- RFC-4502

Fault Recovery

RMON provides the following fault recovery functions.

Interface Down — When an RMON-enabled interface goes down, monitoring continues. However, all data values are registered as 0xFFFFFFFF (32 bits) or 0xFFFFFFFFFFFFFFFF (64 bits). When the interface comes back up, RMON monitoring processes resumes.

 **NOTE:** A network management system (NMS) should be ready to interpret a down interface and plot the interface performance graph accordingly.

Setting the rmon Alarm

To set an alarm on any MIB object, use the `rmon alarm` or `rmon hc-alarm` command in GLOBAL CONFIGURATION mode.

- Set an alarm on any MIB object.
CONFIGURATION mode

```
[no] rmon alarm number variable interval {delta | absolute} rising-threshold  
[value event-number] falling-threshold value event-number [owner string]
```

OR

```
[no] rmon hc-alarm number variable interval {delta | absolute} rising-  
threshold value event-number falling-threshold value event-number [owner  
string]
```

Configure the alarm using the following optional parameters:

- *number*: alarm number, an integer from 1 to 65,535, the value must be unique in the RMON Alarm Table.
- *variable*: the MIB object to monitor — the variable must be in SNMP OID format; for example, 1.3.6.1.2.1.1.3. The object type must be a 32-bit integer for the `rmon alarm` command and 64 bits for the `rmon hc-alarm` command.
- *interval*: time in seconds the alarm monitors the MIB variable, the value must be between 1 to 3,600.
- *delta*: tests the change between MIB variables, this option is the `alarmSampleType` in the RMON Alarm table.
- *absolute*: tests each MIB variable directly, this option is the `alarmSampleType` in the RMON Alarm table.
- *rising-threshold value*: value at which the rising-threshold alarm is triggered or reset. For the `rmon alarm` command, this setting is a 32-bits value, for the `rmon hc-alarm` command, this setting is a 64-bits value.
- *event-number*: event number to trigger when the rising threshold exceeds its limit. This value is identical to the `alarmRisingEventIndex` in the `alarmTable` of the RMON MIB. If there is no corresponding rising-threshold event, the value should be zero.
- *falling-threshold value*: value at which the falling-threshold alarm is triggered or reset. For the `rmon alarm` command, this setting is a 32-bits value, for the `rmon hc-alarm` command this setting is a 64 bits value.

- *event-number*: event number to trigger when the falling threshold exceeds its limit. This value is identical to the alarmFallingEventIndex in the alarmTable of the RMON MIB. If there is no corresponding falling-threshold event, the value should be zero.
- *owner string*: (Optional) specifies an owner for the alarm, this setting is the alarmOwner object in the alarmTable of the RMON MIB. Default is a **null-terminated string**.

Example of the `rmon alarm` Command

To disable the alarm, use the `no` form of the command.

The following example configures RMON alarm number 10. The alarm monitors the MIB variable 1.3.6.1.2.1.2.2.1.20.1 (ifEntry.ifOutErrors) once every 20 seconds until the alarm is disabled, and checks the rise or fall of the variable. The alarm is triggered when the 1.3.6.1.2.1.2.2.1.20.1 value shows a MIB counter increase of 15 or more (such as from 100000 to 100015). The alarm then triggers event number 1, which is configured with the RMON event command. Possible events include a log entry or an SNMP trap. If the 1.3.6.1.2.1.2.2.1.20.1 value changes to 0 (falling-threshold 0), the alarm is reset and can be triggered again.

```
Dell(conf)#rmon alarm 10 1.3.6.1.2.1.2.2.1.20.1 20 delta rising-threshold 15 1
falling-threshold 0
owner nms1
```

Configuring an RMON Event

To add an event in the RMON event table, use the `rmon event` command in GLOBAL CONFIGURATION mode.

- Add an event in the RMON event table.
CONFIGURATION mode

```
[no] rmon event number [log] [trap community] [description string] [owner string]
```

- *number*: assigned event number, which is identical to the eventIndex in the eventTable in the RMON MIB. The value must be an integer from 1 to 65,535 and be unique in the RMON Event Table.
- *log*: (Optional) generates an RMON log entry when the event is triggered and sets the eventType in the RMON MIB to log or log-and-trap. Default is **no log**.
- *trap community*: (Optional) SNMP community string used for this trap. Configures the setting of the eventType in the RMON MIB for this row as either snmp-trap or log-and-trap. This value is identical to the eventCommunityValue in the eventTable in the RMON MIB. Default is `public`.
- *description string*: (Optional) specifies a description of the event, which is identical to the event description in the eventTable of the RMON MIB. The default is a **null-terminated string**.
- *owner string*: (Optional) owner of this event, which is identical to the eventOwner in the eventTable of the RMON MIB. Default is a **null-terminated string**.

Example of the `rmon event` Command

To disable RMON on the interface, use the `no` form of this command.

In the following example, the configuration creates RMON event number 1, with the description "High ifOutErrors", and generates a log entry when an alarm triggers the event. The user `nms1` owns the row that is created in the event table by this command. This configuration also generates an SNMP trap when the event is triggered using the SNMP community string "eventtrap".

```
Dell(conf)#rmon event 1 log trap eventtrap description "High ifOutErrors" owner
nms1
```

Configuring RMON Collection Statistics

To enable RMON MIB statistics collection on an interface, use the `rmon collection statistics` command in INTERFACE CONFIGURATION mode.

- Enable RMON MIB statistics collection.
CONFIGURATION INTERFACE (config-if) mode

```
[no] rmon collection statistics {controlEntry integer} [owner owner-string]
```

- `controlEntry`: specifies the RMON group of statistics using a value.
- `integer`: a value from 1 to 65,535 that identifies the RMON Statistics Table. The value must be unique in the RMON Statistic Table.
- `owner`: (Optional) specifies the name of the owner of the RMON group of statistics.
- `owner-string`: (Optional) records the name of the owner of the RMON group of statistics. The default is a **null-terminated string**.

Example of the `rmon collection statistics` Command

To remove a specified RMON statistics collection, use the `no` form of this command.

The following command example enables the RMON statistics collection on the interface, with an ID value of 20 and an owner of *john*.

```
Dell(conf-if-te-0/40)#rmon collection statistics controlEntry 20 owner john
```

Configuring the RMON Collection History

To enable the RMON MIB history group of statistics collection on an interface, use the `rmon collection history` command in INTERFACE CONFIGURATION mode.

- Configure the RMON MIB history group of statistics collection.
CONFIGURATION INTERFACE (config-if) mode

```
[no] rmon collection history {controlEntry integer} [owner owner-string]  
[buckets bucket-number] [interval seconds]
```

- `controlEntry`: specifies the RMON group of statistics using a value.
- `integer`: a value from 1 to 65,535 that identifies the RMON group of statistics. The value must be a unique index in the RMON History Table.
- `owner`: (Optional) specifies the name of the owner of the RMON group of statistics. The default is a **null-terminated string**.
- `owner-string`: (Optional) records the name of the owner of the RMON group of statistics.
- `buckets`: (Optional) specifies the maximum number of buckets desired for the RMON collection history group of statistics.
- `bucket-number`: (Optional) a value associated with the number of buckets specified for the RMON collection history group of statistics. The value is limited to from 1 to 1000. The default is **50** (as defined in RFC-2819).
- `interval`: (Optional) specifies the number of seconds in each polling cycle.

- `seconds`: (Optional) the number of seconds in each polling cycle. The value is ranged from 5 to 3,600 (Seconds). The default is **1,800** (as defined in RFC-2819).

Example of the `rmon collection history` Command

To remove a specified RMON history group of statistics collection, use the `no` form of this command.

The following command example enables an RMON MIB collection history group of statistics with an ID number of 20 and an owner of *john*, both the sampling interval and the number of buckets use their respective defaults.

```
Dell(conf-if-mgmt)#rmon collection history controlEntry 20 owner john
```

Enabling an RMON MIB Collection History Group

The `rmon collection history` command enables an RMON MIB collection history group of statistics.

In the following example, the command enables an RMON MIB collection history group of statistics with an ID number of 20 and an owner of "john", both the sampling interval and the number of buckets use their respective defaults.

Example of the `rmon collection history` Command

```
Dell(conf-if-te-0/40)#rmon collection history controlEntry 20 owner  
john
```

Rapid Spanning Tree Protocol (RSTP)

Rapid spanning tree protocol (RSTP) is supported on the MXL switch platform.

Protocol Overview

RSTP is a Layer 2 protocol — specified by IEEE 802.1w — that is essentially the same as spanning-tree protocol (STP) but provides faster convergence and interoperability with switches configured with STP and multiple spanning tree protocol (MSTP).

The Dell operating system (OS) supports three other variations of spanning tree, as shown in the following table.

Table 54. Spanning Tree Variations Dell Networking OS Supports

Dell Networking Term	IEEE Specification
Spanning Tree Protocol (STP)	802.1d
Rapid Spanning Tree Protocol (RSTP)	802.1w
Multiple Spanning Tree Protocol (MSTP)	802.1s
Per-VLAN Spanning Tree Plus (PVST+)	Third Party

Configuring Rapid Spanning Tree

Configuring RSTP is a two-step process.

1. Configure interfaces for Layer 2.
2. Enable the rapid spanning tree protocol.

Related Configuration Tasks

- [Adding and Removing Interfaces](#)
- [Modifying Global Parameters](#)
- [Enable BPDU Filtering Globally](#)
- [Modifying Interface Parameters](#)
- [Configuring an EdgePort](#)
- [Prevent Network Disruptions with BPDU Guard](#)
- [Influencing RSTP Root Selection](#)
- [SNMP Traps for Root Elections and Topology Changes](#)
- [Configure Spanning Tree](#)
- [Configuring Fast Hellos for Link State Detection](#)

- [Flush MAC Addresses after a Topology Change](#)

Important Points to Remember

- RSTP is disabled by default.
- The Dell Networking OS supports only one Rapid Spanning Tree (RST) instance.
- All interfaces in virtual local area networks (VLANs) and all enabled interfaces in Layer 2 mode are automatically added to the RST topology.
- Adding a group of ports to a range of VLANs sends multiple messages to the rapid spanning tree protocol (RSTP) task, avoid using the `range` command. When using the `range` command, Dell Networking recommends limiting the range to five ports and 40 VLANs.

Configuring Interfaces for Layer 2 Mode

To configure and enable interfaces in Layer 2 mode, use the following commands. All interfaces on all bridges that participate in Rapid Spanning Tree must be in Layer 2 and enabled.

1. If the interface has been assigned an IP address, remove it.

```
INTERFACE mode
```

```
no ip address
```

2. Place the interface in Layer 2 mode.

```
INTERFACE mode
```

```
switchport
```

3. Enable the interface.

```
INTERFACE mode
```

```
no shutdown
```

Example of Verifying an Interface is in Layer 2 Mode and Enabled

To verify that an interface is in Layer 2 mode and enabled, use the `show config` command from INTERFACE mode. The bold lines indicate that the interface is in Layer 2 mode.

```
Dell(conf-if-te-1/1)#show config
!  
interface TenGigabitEthernet 1/1  
no ip address  
switchport  
no shutdown  
Dell(conf-if-te-1/1)#
```

Enabling Rapid Spanning Tree Protocol Globally

Enable RSTP globally on all participating bridges; it is not enabled by default.

When you enable RSTP, all physical and port-channel interfaces that are enabled and in Layer 2 mode are automatically part of the RST topology.

- Only one path from any bridge to any other bridge is enabled.
- Bridges block a redundant path by disabling one of the link ports.

To enable RSTP globally for all Layer 2 interfaces, use the following commands.

1. Enter PROTOCOL SPANNING TREE RSTP mode.
CONFIGURATION mode

```
protocol spanning-tree rstp
```

2. Enable RSTP.
PROTOCOL SPANNING TREE RSTP mode

```
no disable
```

Example of Verifying that RSTP is Enabled

Example of the `show spanning-tree rstp` Command

Example of the `show spanning-tree rstp brief` Command

To disable RSTP globally for all Layer 2 interfaces, enter the `disable` command from PROTOCOL SPANNING TREE RSTP mode.

To verify that RSTP is enabled, use the `show config` command from PROTOCOL SPANNING TREE RSTP mode. The bold line indicates that RSTP is enabled.

```
Dell(conf-rstp)#show config
!  
protocol spanning-tree rstp  
no disable  
Dell(conf-rstp)#
```





Figure 108. Rapid Spanning Tree Enabled Globally

To view the interfaces participating in RSTP, use the `show spanning-tree rstp` command from EXEC privilege mode. If a physical interface is part of a port channel, only the port channel is listed in the command output.

```
Dell#show spanning-tree rstp
Root Identifier has priority 32768, Address 0001.e801.cbb4
Root Bridge hello time 2, max age 20, forward delay 15, max hops 0
Bridge Identifier has priority 32768, Address 0001.e801.cbb4
Configured hello time 2, max age 20, forward delay 15, max hops 0
We are the root
Current root has priority 32768, Address 0001.e801.cbb4
Number of topology changes 4, last change occurred 00:02:17 ago on TenGig 1/26

Port 377 (TenGigabitethernet 2/1) is designated Forwarding
Port path cost 20000, Port priority 128, Port Identifier 128.377
Designated root has priority 32768, address 0001.e801.cbb4
Designated bridge has priority 32768, address 0001.e801.cbb4
Designated port id is 128.377, designated path cost 0
Number of transitions to forwarding state 1
BPDU : sent 121, received 9
The port is not in the Edge port mode, bpdu filter is disabled

Port 378 (TenGigabitethernet 2/2) is designated Forwarding
Port path cost 20000, Port priority 128, Port Identifier 128.378
Designated root has priority 32768, address 0001.e801.cbb4
Designated bridge has priority 32768, address 0001.e801.cbb4
Designated port id is 128.378, designated path cost 0
Number of transitions to forwarding state 1
```

```
BPDU : sent 121, received 2
The port is not in the Edge port mode, bpdu filter is disabled
```

```
Port 379 (TenGigabitEthernet 2/3) is designated Forwarding
Port path cost 20000, Port priority 128, Port Identifier 128.379
Designated root has priority 32768, address 0001.e801.cbb4
Designated bridge has priority 32768, address 0001.e801.cbb4
Designated port id is 128.379, designated path cost 0
Number of transitions to forwarding state 1
BPDU : sent 121, received 5
The port is not in the Edge port mode, bpdu filter is disabled
```

```
Port 380 (TenGigabitEthernet 2/4) is designated Forwarding
Port path cost 20000, Port priority 128, Port Identifier 128.380
Designated root has priority 32768, address 0001.e801.cbb4
Designated bridge has priority 32768, address 0001.e801.cbb4
Designated port id is 128.380, designated path cost 0
Number of transitions to forwarding state 1
BPDU : sent 147, received 3
The port is not in the Edge port mode, bpdu filter is disabled
```

```
Dell#
```

To confirm that a port is participating in RSTP, use the `show spanning-tree rstp brief` command from EXEC privilege mode.

```
R3#show spanning-tree rstp brief
Executing IEEE compatible Spanning Tree Protocol
Root ID Priority 32768, Address 0001.e801.cbb4
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID Priority 32768, Address 0001.e80f.1dad
Configured hello time 2, max age 20, forward delay 15
Interface                               Designated
Name      PortID  Prio Cost   Sts Cost  Bridge ID          PortID
-----
TenGig 3/1 128.681 128 20000  BLK 20000 32768 0001.e80b.88bd 128.469
TenGig 3/2 128.682 128 20000  BLK 20000 32768 0001.e80b.88bd 128.470
TenGig 3/3 128.683 128 20000  FWD 20000 32768 0001.e801.cbb4 128.379
TenGig 3/4 128.684 128 20000  BLK 20000 32768 0001.e801.cbb4 128.380
Interface
Name      Role  PortID  Prio Cost   Sts Cost  Link-type Edge Bpdu
-----
TenGig 3/1 Altr 128.681 128 20000  BLK 20000  P2P      No  No
TenGig 3/2 Altr 128.682 128 20000  BLK 20000  P2P      No  No
TenGig 3/3 Root 128.683 128 20000  FWD 20000  P2P      No  No
TenGig 3/4 Altr 128.684 128 20000  BLK 20000  P2P      No  No
R3#
```

Adding and Removing Interfaces

To add and remove interfaces, use the following commands.

To add an interface to the Rapid Spanning Tree topology, configure it for Layer 2 and it is automatically added. If you previously disabled RSTP on the interface using the `no spanning-tree 0` command, re-enable it using the `spanning-tree 0` command.

- Remove an interface from the Rapid Spanning Tree topology.
`no spanning-tree 0`


For bridge protocol data units (BPDU) filtering behavior, refer to [Removing an Interface from the Spanning Tree Group](#).

Modifying Global Parameters

You can modify RSTP parameters.

The root bridge sets the values for forward-delay, hello-time, and max-age and overwrites the values set on other bridges participating in the Rapid Spanning Tree group.

- **Forward-delay** — the amount of time an interface waits in the Listening state and the Learning state before it transitions to the Forwarding state.
- **Hello-time** — the time interval in which the bridge sends RSTP BPDUs.
- **Max-age** — the length of time the bridge maintains configuration information before it refreshes that information by recomputing the RST topology.

 **NOTE:** Dell Networking recommends that only experienced network administrators change the Rapid Spanning Tree group parameters. Poorly planned modification of the RSTP parameters can negatively affect network performance.

The following table displays the default values for RSTP.

Table 55. RSTP Default Values

RSTP Parameter	Default Value
Forward Delay	15 seconds
Hello Time	2 seconds
Max Age	20 seconds
Port Cost:	Port Cost:
• 10-Gigabit Ethernet interfaces	• 1400
• 40-Gigabit Ethernet interfaces	• 2000
• Port Channel with two 40-Gigabit Ethernet interfaces	• 600
• Port Channel with two 10-Gigabit Ethernet interfaces	• 1800
Port Priority	128

To change these parameters, use the following commands.

- Change the forward-delay parameter.
PROTOCOL SPANNING TREE RSTP mode


```
forward-delay seconds
```

The range is from 4 to 30.

The default is **15 seconds**.

- Change the hello-time parameter.
PROTOCOL SPANNING TREE RSTP mode

```
hello-time seconds
```

 **NOTE:** With large configurations (especially those configurations with more ports) Dell Networking recommends increasing the hello-time.

The range is from 1 to 10.

The default is **2 seconds**.

- Change the max-age parameter.

PROTOCOL SPANNING TREE RSTP mode

`max-age seconds`

The range is from 6 to 40.

The default is **20 seconds**.

To view the current values for global parameters, use the `show spanning-tree rstp` command from EXEC privilege mode.

Enable BPDU Filtering Globally

The enabling of BPDU Filtering stops transmitting of BPDUs on the operational port fast enabled ports by default.

When BPDUs are received, the spanning tree is automatically prepared. By default global bpdu filtering is disabled.

Enable BPDU Filter globally to filter transmission of BPDU port fast enabled interfaces.

PROTOCOL SPANNING TREE RSTP mode

edge-port bpdu filter default

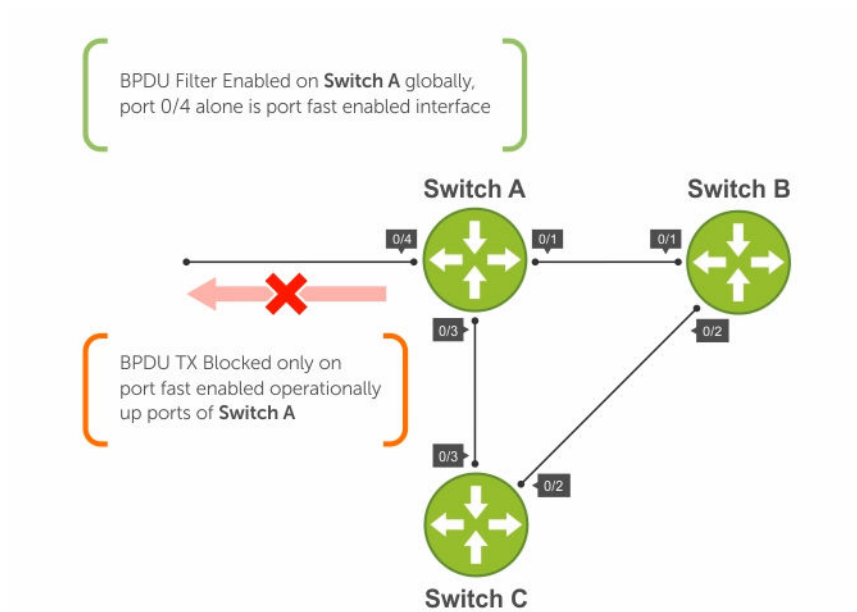


Figure 109. BPDUs Filtering Enabled Globally

Modifying Interface Parameters

On interfaces in Layer 2 mode, you can set the port cost and port priority values.

- **Port cost** — a value that is based on the interface type. The previous table lists the default values. The greater the port cost, the less likely the port is selected to be a forwarding port.
- **Port priority** — influences the likelihood that a port is selected to be a forwarding port in case that several ports have the same port cost.

To change the port cost or priority of an interface, use the following commands.

- Change the port cost of an interface.

INTERFACE mode

```
spanning-tree rstp cost cost
```

The range is from 0 to 65535.

The default is listed in the previous table.

- Change the port priority of an interface.

INTERFACE mode

```
spanning-tree rstp priority priority-value
```


The range is from 0 to 240.

The default is **128**.

To view the current values for interface parameters, use the `show spanning-tree rstp` command from EXEC privilege mode.

Configuring an EdgePort

The EdgePort feature enables interfaces to begin forwarding traffic approximately 30 seconds sooner. In this mode an interface forwards frames by default until it receives a BPDU that indicates that it should behave otherwise; it does not go through the Learning and Listening states. The `bpduguard shutdown-on-violation` option causes the interface hardware to be shut down when it receives a BPDU. When only `bpduguard` is implemented, although the interface is placed in an Error Disabled state when receiving the BPDU, the physical interface remains up and spanning-tree drops packets in the hardware after a BPDU violation. BPDUs are dropped in the software after receiving the BPDU violation. This feature is the same as PortFast mode in Spanning Tree.

 **CAUTION: Configure EdgePort only on links connecting to an end station. If you enable EdgePort on an interface connected to a network, it can cause loops.**

Dell Networking OS Behavior: Regarding `bpduguard shutdown-on-violation` behavior:

- If the interface to be shut down is a port channel, all the member ports are disabled in the hardware.
- When you add a physical port to a port channel already in the Error Disable state, the new member port is also disabled in the hardware.
- When you remove a physical port from a port channel in the Error Disable state, the error disabled state is cleared on this physical port (the physical port is enabled in the hardware).
- You can clear the Error Disabled state with any of the following methods:
 - Perform an `shutdown` command on the interface.
 - Disable the `shutdown-on-violation` command on the interface (the `no spanning-tree stp-id portfast [bpduguard | [shutdown-on-violation]]` command).
 - Disable spanning tree on the interface (the `no spanning-tree` command in INTERFACE mode).
 - Disable global spanning tree (the `no spanning-tree` command in CONFIGURATION mode).


To enable EdgePort on an interface, use the following command.

- Enable EdgePort on an interface.
INTERFACE mode

```
spanning-tree rstp edge-port [bpduguard | shutdown-on-violation][bpdufilter]
```

Example of Verifying an EdgePort is Enabled on an Interface

To verify that EdgePort is enabled on a port, use the `show spanning-tree rstp` command from EXEC privilege mode or the `show config` command from INTERFACE mode.

 **NOTE:** Dell Networking recommends using the `show config` command from INTERFACE mode.

In the following example, the bold line indicates that the interface is in EdgePort mode.

```
Dell(conf-if-te-2/0)#show config
!
interface TenGigabitEthernet 2/0
  no ip address
  switchport
  spanning-tree rstp edge-port
```

```
shutdown
Dell(conf-if-te-2/0)#
```

Influencing RSTP Root Selection

RSTP determines the root bridge, but you can assign one bridge a lower priority to increase the likelihood that it is selected as the root bridge.

To change the bridge priority, use the following command.

- Assign a number as the bridge priority or designate it as the primary or secondary root.
PROTOCOL SPANNING TREE RSTP mode

```
bridge-priority priority-value
```

- *priority-value* The range is from 0 to 65535. The lower the number assigned, the more likely this bridge becomes the root bridge.

The default is **32768**. Entries must be multiples of 4096.

Example of the bridge-priority Command

A console message appears when a new root bridge has been assigned. The following example shows the console message after the `bridge-priority` command is used to make R2 the root bridge (shown in bold).

```
Dell(conf-rstp)#bridge-priority 4096
04:27:59: %RPM0-P:RP2 %SPANMGR-5-STP_ROOT_CHANGE: RSTP root changed. My Bridge
ID:
4096:0001.e80b.88bd Old Root: 32768:0001.e801.cbb4 New Root: 4096:0001.e80b.88bd
```

SNMP Traps for Root Elections and Topology Changes

To enable SNMP traps for RSTP, MSTP, and PVST+ collectively, use the following command.

Enable SNMP traps for RSTP, MSTP, and PVST+ collectively.

```
snmp-server enable traps xstp
```

Configuring Fast Hellos for Link State Detection

To achieve sub-second link-down detection so that convergence is triggered faster, use RSTP fast hellos. The standard RSTP link-state detection mechanism does not offer the same low link-state detection speed.

RSTP fast hellos decrease the hello interval to the order of milliseconds and all timers derived from the hello timer are adjusted accordingly. This feature does not inter-operate with other vendors, and is available only for RSTP.


- Configure a hello time on the order of milliseconds.
PROTOCOL RSTP mode

```
hello-time milli-second interval
```

The range is from 50 to 950 milliseconds.

Example of Verifying Hello-Time Interval

```
Dell(conf-rstp)#do show spanning-tree rstp brief
Executing IEEE compatible Spanning Tree Protocol
Root ID      Priority 0, Address 0001.e811.2233
Root Bridge hello time 50 ms, max age 20, forward delay 15
Bridge ID    Priority 0, Address 0001.e811.2233
We are the root
Configured hello time 50 ms, max age 20, forward delay 15
```

 **NOTE:** The hello time is encoded in BPDUs in increments of 1/256ths of a second. The standard minimum hello time in seconds is 1 second, which is encoded as 256. Millisecond hello times are encoded using values less than 256; the millisecond hello time equals $(x/1000)*256$. When you configure millisecond hellos, the default hello interval of 2 seconds is still used for edge ports; the millisecond hello interval is not used.

Security

Security features are supported on the MXL switch platform.

This chapter describes several ways to provide access security to the Dell Networking system.

For details about all the commands described in this chapter, refer to the *Security* chapter in the *Dell Networking OS Command Reference Guide*.

AAA Accounting

Accounting, authentication, and authorization (AAA) accounting is part of the AAA security model.

For details about commands related to AAA security, refer to the *Security* chapter in the *Dell Networking OS Command Reference Guide*.

AAA accounting enables tracking of services that users are accessing and the amount of network resources being consumed by those services. When you enable AAA accounting, the network server reports user activity to the security server in the form of accounting records. Each accounting record is comprised of accounting attribute/value (AV) pairs and is stored on the access control server.

As with authentication and authorization, you must configure AAA accounting by defining a named list of accounting methods and then applying that list to various interfaces.

Configuration Task List for AAA Accounting

The following sections present the AAA accounting configuration tasks.

- [Enabling AAA Accounting](#) (mandatory)
- [Suppressing AAA Accounting for Null Username Sessions](#) (optional)
- [Configuring Accounting of EXEC and Privilege-Level Command Usage](#) (optional)
- [Configuring AAA Accounting for Terminal Lines](#) (optional)
- [Monitoring AAA Accounting](#) (optional)

Enabling AAA Accounting

The `aaa accounting` command allows you to create a record for any or all of the accounting functions monitored.

To enable AAA accounting, use the following command.

- Enable AAA accounting and create a record for monitoring the accounting function.
CONFIGURATION mode

```
aaa accounting {commands | exec | suppress | system} {default | name} {start-stop | wait-start | stop-only} {tacacs+}
```

The variables are:

- `command level`: sends accounting of commands executed at the specified privilege level.
- `exec`: sends accounting information when a user has logged in to EXEC mode.
- `suppress`: do not generate accounting records for a specific type of user.
- `system`: sends accounting information of any other AAA configuration.
- `default | name`: enter the name of a list of accounting methods.
- `start-stop`: use for more accounting information, to send a start-accounting notice at the beginning of the requested event and a stop-accounting notice at the end.
- `wait-start`: ensures that the TACACS+ security server acknowledges the start notice before granting the user's process request.
- `stop-only`: use for minimal accounting; instructs the TACACS+ server to send a stop record accounting notice at the end of the requested user process.
- `tacacs+`: designate the security service. Currently, the system supports only TACACS+.

Suppressing AAA Accounting for Null Username Sessions

When you activate AAA accounting, the Dell Networking OS issues accounting records for all users on the system, including users whose username string, because of protocol translation, is NULL. An example of this is a user who comes in on a line where the AAA authentication `login method-list none` command is applied. To prevent accounting records from being generated for sessions that do not have usernames associated with them, use the following command.

- Prevent accounting records from being generated for users whose username string is NULL.
CONFIGURATION mode

```
aaa accounting suppress null-username
```

Configuring Accounting of EXEC and Privilege-Level Command Usage

The network access server monitors the accounting functions defined in the TACACS+ attribute/value (AV) pairs.

- Configure AAA accounting to monitor accounting functions defined in TACACS+.
CONFIGURATION mode

```
aaa accounting system default start-stop tacacs+
```

```
aaa accounting command 15 default start-stop tacacs+
```

System accounting can use only the default method list.

Example of Configuring AAA Accounting to Track EXEC and EXEC Privilege Level Command Use

In the following sample configuration, AAA accounting is set to track all usage of EXEC commands and commands on privilege level 15.

```
Dell(conf)#aaa accounting exec default start-stop tacacs+  
Dell(conf)#aaa accounting command 15 default start-stop tacacs+
```

Configuring AAA Accounting for Terminal Lines

To enable AAA accounting with a named method list for a specific terminal line (where *com15* and *execAcct* are the method list names), use the following commands.

- Configure AAA accounting for terminal lines.

CONFIG-LINE-VTY mode

```
accounting commands 15 com15
```

```
accounting exec execAcct
```

Example of Enabling AAA Accounting with a Named Method List

```
Dell(config-line-vty)# accounting commands 15 com15
```

```
Dell(config-line-vty)# accounting exec execAcct
```

Monitoring AAA Accounting

The Dell Networking OS does not support periodic interim accounting because the *periodic* command can cause heavy congestion when many users are logged in to the network.

No specific *show* command exists for TACACS+ accounting.

To obtain accounting records displaying information about users currently logged in, use the following command.

- Step through all active sessions and print all the accounting records for the actively accounted functions.

CONFIGURATION mode or EXEC Privilege mode

```
show accounting
```

Example of the *show accounting* Command for AAA Accounting

```
Dell#show accounting
```

```
Active accounted actions on tty2, User admin Priv 1
```

```
Task ID 1, EXEC Accounting record, 00:00:39 Elapsed, service=shell
```

```
Active accounted actions on tty3, User admin Priv 1
```

```
Task ID 2, EXEC Accounting record, 00:00:26 Elapsed, service=shell
```

```
Dell#
```


AAA Authentication

The Dell Networking OS supports a distributed client/server system implemented through authentication, authorization, and accounting (AAA) to help secure networks against unauthorized access.

In the Dell Networking implementation, the Dell Networking system acts as a RADIUS or TACACS+ client and sends authentication requests to a central remote authentication dial-in service (RADIUS) or Terminal access controller access control system plus (TACACS+) server that contains all user authentication and network service access information.

Dell Networking uses local usernames/passwords (stored on the Dell Networking system) or AAA for login authentication. With AAA, you can specify the security protocol or mechanism for different login methods

and different users. In the Dell Networking OS, AAA uses a list of authentication methods, called method lists, to define the types of authentication and the sequence in which they are applied. You can define a method list or use the default method list. User-defined method lists take precedence over the default method list.

 **NOTE:** If a console user logs in with RADIUS authentication, the privilege level is applied from the RADIUS server if the privilege level is configured for that user in RADIUS, whether you configure RADIUS authorization.

Configuration Task List for AAA Authentication

The following sections provide the configuration tasks.

- [Configure Login Authentication for Terminal Lines](#)
- [Configuring AAA Authentication Login Methods](#)
- [Enabling AAA Authentication](#)
- [Enabling AAA Authentication – RADIUS](#)

For a complete list of all commands related to login authentication, refer to the *Security* chapter in the *Dell Networking OS Command Reference Guide*.

Configure Login Authentication for Terminal Lines

You can assign up to five authentication methods to a method list. The Dell Networking OS evaluates the methods in the order in which you enter them in each list.

If the first method list does not respond or returns an error, the system applies the next method list until the user either passes or fails the authentication. If the user fails a method list, the system does not apply the next method list.

Configuring AAA Authentication Login Methods

To configure an authentication method and method list, use the following commands.

Dell Networking OS Behavior: If you use a method list on the console port in which RADIUS or TACACS is the last authentication method, and the server is not reachable, the system allows access even though the username and password credentials cannot be verified. Only the console port behaves this way, and does so to ensure that users are not locked out of the system if network-wide issue prevents access to these servers.

1. Define an authentication method-list (*method-list-name*) or specify the default.

CONFIGURATION mode

```
aaa authentication login {method-list-name | default} method1 [... method4]
```

The default method-list is applied to all terminal lines.

Possible methods are:

- `enable`: use the password you defined using the `enable secret` or `enable password` command in CONFIGURATION mode.
- `line`: use the password you defined using the `password` command in LINE mode.
- `local`: use the username/password database defined in the local configuration.

- none: no authentication.
 - radius: use the RADIUS servers configured with the radius-server host command.
 - tacacs+: use the TACACS+ servers configured with the tacacs-server host command.
2. Enter LINE mode.

CONFIGURATION mode


```
line {aux 0 | console 0 | vty number [... end-number]}
```

3. Assign a *method-list-name* or the default list to the terminal line.

LINE mode

```
login authentication {method-list-name | default}
```

To view the configuration, use the `show config` command in LINE mode or the `show running-config` in EXEC Privilege mode.

 **NOTE:** Dell Networking recommends using the `none` method only as a backup. This method does not authenticate users. The `none` and `enable` methods do not work with secure shell (SSH).

You can create multiple method lists and assign them to different terminal lines.

Enabling AAA Authentication

To enable AAA authentication, use the following command.

- Enable AAA authentication.

CONFIGURATION mode

```
aaa authentication enable {method-list-name | default} method1 [... method4]
```

- `default`: uses the listed authentication methods that follow this argument as the default list of methods when a user logs in.
- `method-list-name`: character string used to name the list of enable authentication methods activated when a user logs in.
- `method1 [... method4]`: any of the following: RADIUS, TACACS, enable, line, none.

If you do not set the default list, only the local enable is checked. This setting has the same effect as issuing an `aaa authentication enable default enable` command.

Enabling AAA Authentication — RADIUS

To enable authentication from the RADIUS server, and use TACACS as a backup, use the following commands.

1. Enable RADIUS and set up TACACS as backup.

CONFIGURATION mode

```
aaa authentication enable default radius tacacs
```

2. Establish a host address and password.

CONFIGURATION mode

```
radius-server host x.x.x.x key some-password
```

3. Establish a host address and password.

CONFIGURATION mode

```
tacacs-server host x.x.x.x key some-password
```

Example of Enabling Authentication from the RADIUS Server

Example of Enabling Local Authentication for the Console and Remote Authentication for VTY Lines

To get enable authentication from the RADIUS server and use TACACS as a backup, issue the following commands.

```
Dell(config)# aaa authentication enable default radius tacacs
Radius and TACACS server has to be properly setup for this.
Dell(config)# radius-server host x.x.x.x key <some-password>
Dell(config)# tacacs-server host x.x.x.x key <some-password>
```

To use local authentication for enable secret on the console, while using remote authentication on VTY lines, issue the following commands.

```
Dell(config)# aaa authentication enable mymethodlist radius tacacs
Dell(config)# line vty 0 9
Dell(config-line-vty)# enable authentication mymethodlist
```

Server-Side Configuration

- **TACACS+** — When using TACACS+, Dell Networking sends an initial packet with service type SVC_ENABLE, and then sends a second packet with just the password. The TACACS server must have an entry for username \$enable\$.
- **RADIUS** — When using RADIUS authentication, the system sends an authentication packet with the following:

```
Username: $enab15$
Password: <password-entered-by-user>
```

Therefore, the RADIUS server must have an entry for this username.

AAA Authorization

The Dell Networking OS enables AAA new-model by default.

You can set authorization to be either `local` or `remote`. Different combinations of authentication and authorization yield different results. By default, the system sets both to **local**.

Privilege Levels Overview

Limiting access to the system is one method of protecting the system and your network. However, at times, you might need to allow others access to the router and you can limit that access to a subset of commands. In the Dell Networking OS, you can configure a privilege level for users who need limited access to the system.

Every command in the Dell Networking OS is assigned a privilege level of 0, 1, or 15. You can configure up to 16 privilege levels. The Dell Networking OS is pre-configured with three privilege levels and you can configure 13 more. The three pre-configured levels are:

- **Privilege level 1** — is the default level for EXEC mode. At this level, you can interact with the router, for example, view some `show` commands and Telnet and ping to test connectivity, but you cannot configure the router. This level is often called the “user” level. One of the commands available in Privilege level 1 is the `enable` command, which you can use to enter a specific privilege level.
- **Privilege level 0** — contains only the `end`, `enable`, and `disable` commands.
- **Privilege level 15** — the default level for the `enable` command, is the highest level. In this level you can access any command in the Dell Networking OS.

Privilege levels 2 through 14 are not configured and you can customize them for different users and access.

After you configure other privilege levels, enter those levels by adding the level parameter after the `enable` command or by configuring a user name or password that corresponds to the privilege level. For more information about configuring user names, refer to [Configuring a Username and Password](#).

By default, commands are assigned to different privilege levels. You can access those commands only if you have access to that privilege level. For example, to reach the `protocol spanning-tree` command, log in to the router, enter the `enable` command for privilege level 15 (this privilege level is the default level for the command) and then enter CONFIGURATION mode.

You can configure passwords to control access to the box and assign different privilege levels to users. The Dell Networking OS supports the use of passwords when you log in to the system and when you enter the `enable` command. If you move between privilege levels, you are prompted for a password if you move to a higher privilege level.

Configuration Task List for Privilege Levels

The following list has the configuration tasks for privilege levels and passwords.

- [Configuring a Username and Password](#) (mandatory)
- [Configuring the Enable Password Command](#) (mandatory)
- [Configuring Custom Privilege Levels](#) (mandatory)
- [Specifying LINE Mode Password and Privilege](#) (optional)
- [Enabling and Disabling Privilege Levels](#) (optional)

For a complete listing of all commands related to privilege levels and passwords, refer to the *Security* chapter in the *Dell Networking OS Command Reference Guide*.

Configuring a Username and Password

In the Dell Networking OS, you can assign a specific username to limit user access to the system. To configure a username and password, use the following command.

- Assign a user name and password.
CONFIGURATION mode

```
username name [access-class access-list-name] [nopassword | password
[encryption-type] password] [privilege level] [secret]
```

Configure the optional and required parameters:

- `name`: Enter a text string up to 63 characters long.

- `access-class access-list-name`: Restrict access by access-class.
- `nopassword`: Require password for the user to login.
- `encryption-type`: Enter 0 for plain text or 7 for encrypted text.
- `password`: Enter a string. Specify the password for the user.
- `privilege level`: The range is from 0 to 15.
- `secret`: Specify the secret for the user.

To view username, use the `show users` command in EXEC Privilege mode.

Configuring the Enable Password Command

To configure the Dell Networking OS, use the `enable` command to enter EXEC Privilege level 15. After entering the command, the system requests that you enter a password.

Privilege levels are not assigned to passwords, rather passwords are assigned to a privilege level. You can always change a password for any privilege level. To change to a different privilege level, enter the `enable` command, then the privilege level. If you do not enter a privilege level, the default level **15** is assumed.

To configure a password for a specific privilege level, use the following command.

- Configure a password for a privilege level.

CONFIGURATION mode

```
enable password [level level] [encryption-mode] password
```

Configure the optional and required parameters:

- `level level`: Specify a level from 0 to 15. Level 15 includes all levels.
- `encryption-type`: Enter 0 for plain text or 7 for encrypted text.
- `password`: Enter a string.

To change only the password for the `enable` command, configure only the `password` parameter.

To view the configuration for the `enable secret` command, use the `show running-config` command in EXEC Privilege mode.

In custom-configured privilege levels, the `enable` command is always available. No matter what privilege level you entered, you can enter the `enable 15` command to access and configure all CLIs.

Configuring Custom Privilege Levels

In addition to assigning privilege levels to the user, you can configure the privilege levels of commands so that they are visible in different privilege levels.

Within the Dell Networking OS, commands have certain privilege levels. With the `privilege` command, you can change the default level or you can reset their privilege level back to the default. Assign the launch keyword (for example, `configure`) for the keyword's command mode.

To assign commands and passwords to a custom privilege level, use the following commands. You must be in privilege level 15.

1. Assign a user name and password.

CONFIGURATION mode


```
username name [access-class access-list-name] [privilege level] [nopassword  
| password [encryption-type password] [secret]]
```

Configure the optional and required parameters:

- *name*: enter a text string (up to 63 characters).
- *access-class access-list-name*: enter the name of a configured IP ACL.
- *privilege level*: the range is from 0 to 15.
- *nopassword*: do not require the user to enter a password.
- *encryption-type*: enter 0 for plain text or 7 for encrypted text.
- *password*: enter a text string.
- *secret*: specify the secret for the user.

2. Configure a password for privilege level.

CONFIGURATION mode

```
enable password [level level] [encryption-mode] password
```

Configure the optional and required parameters:

- *level level*: specify a level from 0 to 15. Level 15 includes all levels.
- *encryption-type*: enter 0 for plain text or 7 for encrypted text.
- *password*: enter a text string up to 32 characters long.

To change only the password for the `enable` command, configure only the `password` parameter.

3. Configure level and commands for a mode or reset a command's level.

CONFIGURATION mode

```
privilege mode {level level command | reset command}
```

Configure the following required and optional parameters:

- *mode*: enter a keyword for the modes (`exec`, `configure`, `interface`, `line`, `route-map`, or `router`)
- *level level*: the range is from 0 to 15. Levels 0, 1, and 15 are pre-configured. Levels 2 to 14 are available for custom configuration.
- *command*: an Dell CLI keyword (up to five keywords allowed).
- *reset*: return the command to its default privilege mode.

Example of Configuring a Custom Privilege Level

Example of Privilege Level Login and Available Commands

To view the configuration, use the `show running-config` command in EXEC Privilege mode.

The following example shows a configuration to allow a user *john* to view only EXEC mode commands and all `snmp-server` commands. Because the `snmp-server` commands are *enable* level commands and, by default, found in CONFIGURATION mode, also assign the launch command for CONFIGURATION mode, `configure`, to the same privilege level as the `snmp-server` commands.

Line 1: The user *john* is assigned privilege level 8 and assigned a password.

Line 2: All other users are assigned a password to access privilege level 8.

Line 3: The `configure` command is assigned to privilege level 8 because it needs to reach CONFIGURATION mode where the `snmp-server` commands are located.

Line 4: The `snmp-server` commands, in CONFIGURATION mode, are assigned to privilege level 8.

```
Dell(conf)#username john privilege 8 password john
Dell(conf)#enable password level 8 notjohn
Dell(conf)#privilege exec level 8 configure
Dell(conf)#privilege config level 8 snmp-server
Dell(conf)#end
Dell#show running-config
Current Configuration ...
!
hostname FTOS
!
enable password level 8 notjohn
enable password FTOS
!
username admin password 0 admin
username john password 0 john privilege 8
!
```

The following example shows the Telnet session for user *john*. The `show privilege` command output confirms that *john* is in privilege level 8. In EXEC Privilege mode, *john* can access only the commands listed. In CONFIGURATION mode, *john* can access only the `snmp-server` commands.

```
apollo% telnet 172.31.1.53
Trying 172.31.1.53...
Connected to 172.31.1.53.
Escape character is '^]'.
Login: john
Password:
Dell#show priv
Current privilege level is 8
Dell#?
configure          Configuring from terminal
disable            Turn off privileged commands
enable             Turn on privileged commands
exit               Exit from the EXEC
no                Negate a command
show              Show running system information
terminal          Set terminal line parameters
traceroute        Trace route to destination
Dell#confi
Dell(conf)#?
end               Exit from Configuration mode
```

Specifying LINE Mode Password and Privilege

You can specify a password authentication of all users on different terminal lines.

The user's privilege level is the same as the privilege level assigned to the terminal line, unless a more specific privilege level is assigned to the user.

To specify a password for the terminal line, use the following commands.

- Configure a custom privilege level for the terminal lines.
LINE mode

```
privilege level level
```

- *level level*: The range is from 0 to 15. Levels 0, 1, and 15 are pre-configured. Levels 2 to 14 are available for custom configuration.
- Specify either a plain text or encrypted password.
LINE mode

```
password [encryption-type] password
```

Configure the following optional and required parameters:

- *encryption-type*: Enter 0 for plain text or 7 for encrypted text.
- *password*: Enter a text string up to 25 characters long.

To view the password configured for a terminal, use the `show config` command in LINE mode.

Enabling and Disabling Privilege Levels

To enable and disable privilege levels, use the following commands.

- Set a user's security level.
EXEC Privilege mode

```
enable or enable privilege-level
```

If you do not enter a privilege level, the system sets it to 15 by default.

- Move to a lower privilege level.
EXEC Privilege mode

```
disable level-number
```

- *level-number*: The level-number you wish to set.

If you enter `disable` without a level-number, your security level is 1.

RADIUS

Remote authentication dial-in user service (RADIUS) is a distributed client/server protocol.

This protocol transmits authentication, authorization, and configuration information between a central RADIUS server and a RADIUS client (the Dell Networking system). The system sends user information to the RADIUS server and requests authentication of the user and password. The RADIUS server returns one of the following responses:

- **Access-Accept** — the RADIUS server authenticates the user.
- **Access-Reject** — the RADIUS server does not authenticate the user.

If an error occurs in the transmission or reception of RADIUS packets, you can view the error by enabling the `debug radius` command.

Transactions between the RADIUS server and the client are encrypted (the users' passwords are not sent in plain text). RADIUS uses UDP as the transport protocol between the RADIUS server host and the client.

For more information about RADIUS, refer to RFC 2865, *Remote Authentication Dial-in User Service*.

RADIUS Authentication and Authorization


The Dell Networking OS supports RADIUS for user authentication (text password) at login and can be specified as one of the login authentication methods in the `aaa authentication login` command.

When configuring AAA authorization, you can configure to limit the attributes of services available to a user. When you enable authorization, the network access server uses configuration information from the user profile to issue the user's session. The user's access is limited based on the configuration attributes.

RADIUS exec-authorization stores a user-shell profile and that is applied during user login. You may name the relevant named-lists with either a unique name or the default name. When you enable authorization by the RADIUS server, the server returns the following information to the client:

- [Idle Time](#)
- [ACL Configuration Information](#)
- [Auto-Command](#)
- [Privilege Levels Overview](#)

After gaining authorization for the first time, you may configure these attributes.

 **NOTE:** RADIUS authentication/authorization is done for every login. There is no difference between first-time login and subsequent logins.

Idle Time

Every session line has its own idle-time. If the idle-time value is not changed, the default value of **30 minutes** is used.

RADIUS specifies idle-time allow for a user during a session before timeout. When a user logs in, the lower of the two idle-time values (configured or default) is used. The idle-time value is updated if both of the following happens:

- The administrator changes the idle-time of the line on which the user has logged in.
- The idle-time is lower than the RADIUS-returned idle-time.


ACL Configuration Information

The RADIUS server can specify an ACL. If an ACL is configured on the RADIUS server, and if that ACL is present, the user may be allowed access based on that ACL.

If the ACL is absent, authorization fails, and a message is logged indicating this.

RADIUS can specify an ACL for the user if both of the following are true:

- If an ACL is absent.
- If there is a very long delay for an entry, or a denied entry because of an ACL, and a message is logged.

 **NOTE:** The ACL name must be a string. Only standard ACLs in authorization (both RADIUS and TACACS) are supported. Authorization is denied in cases using Extended ACLs.

Auto-Command

You can configure the system through the RADIUS server to automatically execute a command when you connect to a specific line.

The `auto-command` command is executed when the user is authenticated and before the prompt appears to the user.

- Automatically execute a command.

```
auto-command
```

Setting Access to Privilege Levels through RADIUS

To configure a privilege level for the user to enter into when they connect to a session, use the following command.

```
Configure a privilege level for the user to enter into when they connect to a session through the RADIUS server.
```

```
privilege level
```

```
Configure this value on the client system.
```


Configuration Task List for RADIUS

To authenticate users using RADIUS, you must specify at least one RADIUS server so that the system can communicate with and configure RADIUS as one of your authentication methods.

The following list includes the configuration tasks for RADIUS.

- [Defining a AAA Method List to be Used for RADIUS](#) (mandatory)
- [Applying the Method List to Terminal Lines](#) (mandatory except when using default lists)
- [Specifying a RADIUS Server Host](#) (mandatory)
- [Setting Global Communication Parameters for all RADIUS Server Hosts](#) (optional)
- [Monitoring RADIUS](#) (optional)

For a complete listing of all Dell Networking OS commands related to RADIUS, refer to the *Security* chapter in the *Dell Networking OS Command Reference Guide*.

 **NOTE:** RADIUS authentication and authorization are done in a single step. Hence, authorization cannot be used independent of authentication. However, if you have configured RADIUS authorization and have not configured authentication, a message is logged stating this. During authorization, the next method in the list (if present) is used, or if another method is not present, an error is reported.

To view the configuration, use the `show config` in LINE mode or the `show running-config` command in EXEC Privilege mode.

Defining a AAA Method List to be Used for RADIUS

To configure RADIUS to authenticate or authorize users on the system, create a AAA method list. Default method lists do not need to be explicitly applied to the line, so they are not mandatory.

To create a method list, use the following commands.

- Enter a text string (up to 16 characters long) as the name of the method list you wish to use with the RADIUS authentication method.

CONFIGURATION mode

```
aaa authentication login method-list-name radius
```

- Create a method list with RADIUS and TACACS+ as authorization methods.

CONFIGURATION mode

```
aaa authorization exec {method-list-name | default} radius tacacs+
```

Typical order of methods: RADIUS, TACACS+, Local, None.

If RADIUS denies authorization, the session ends (RADIUS must not be the last method specified).

Applying the Method List to Terminal Lines

To enable RADIUS AAA login authentication for a method list, apply it to a terminal line.

To configure a terminal line for RADIUS authentication and authorization, use the following commands.

- Enter LINE mode.

CONFIGURATION mode

```
line {aux 0 | console 0 | vty number [end-number]}
```

- Enable AAA login authentication for the specified RADIUS method list.

LINE mode

```
login authentication {method-list-name | default}
```

This procedure is mandatory if you are not using default lists.

- To use the method list.

CONFIGURATION mode

```
authorization exec methodlist
```

Specifying a RADIUS Server Host

When configuring a RADIUS server host, you can set different communication parameters, such as the UDP port, the key password, the number of retries, and the timeout.

To specify a RADIUS server host and configure its communication parameters, use the following command.

- Enter the host name or IP address of the RADIUS server host.

CONFIGURATION mode

```
radius-server host {hostname | ip-address} [auth-port port-number]  
[retransmit retries] [timeout seconds] [key [encryption-type] key]
```

Configure the optional communication parameters for the specific host:

- *auth-port* *port-number*: the range is from 0 to 65535. Enter a UDP port number. The default is **1812**.
- *retransmit* *retries*: the range is from 0 to 100. Default is **3**.

- `timeout seconds`: the range is from 0 to 1000. Default is **5 seconds**.
- `key [encryption-type] key`: enter 0 for plain text or 7 for encrypted text, and a string for the key. The key can be up to 42 characters long. This key must match the key configured on the RADIUS server host.

If you do not configure these optional parameters, the global default values for all RADIUS host are applied.

To specify multiple RADIUS server hosts, configure the `radius-server host` command multiple times. If you configure multiple RADIUS server hosts, the system attempts to connect with them in the order in which they were configured. When the system attempts to authenticate a user, the software connects with the RADIUS server hosts one at a time, until a RADIUS server host responds with an accept or reject response.

If you want to change an optional parameter setting for a specific host, use the `radius-server host` command. To change the global communication settings to all RADIUS server hosts, refer to [Setting Global Communication Parameters for all RADIUS Server Hosts](#).

To view the RADIUS configuration, use the `show running-config radius` command in EXEC Privilege mode.

To delete a RADIUS server host, use the `no radius-server host {hostname | ip-address}` command.

Setting Global Communication Parameters for all RADIUS Server Hosts

You can configure global communication parameters (`auth-port`, `key`, `retransmit`, and `timeout` parameters) and specific host communication parameters on the same system. However, if you configure both global and specific host parameters, the specific host parameters override the global parameters for that RADIUS server host.

To set global communication parameters for all RADIUS server hosts, use the following commands.

- Set a time interval after which a RADIUS host server is declared dead.
CONFIGURATION mode

```
radius-server deadtime seconds
```

- `seconds`: the range is from 0 to 2147483647. The default is **0 seconds**.

- Configure a key for all RADIUS communications between the system and RADIUS server hosts.
CONFIGURATION mode

```
radius-server key [encryption-type] key
```

- `encryption-type`: enter 7 to encrypt the password. Enter 0 to keep the password as plain text.
- `key`: enter a string. The key can be up to 42 characters long. You cannot use spaces in the key.

- Configure the number of times the system retransmits RADIUS requests.
CONFIGURATION mode

```
radius-server retransmit retries
```

- `retries`: the range is from 0 to 100. Default is **3 retries**.

- Configure the time interval the system waits for a RADIUS server host response.
CONFIGURATION mode

```
radius-server timeout seconds
```

- *seconds*: the range is from 0 to 1000. Default is **5 seconds**.

To view the configuration of RADIUS communication parameters, use the `show running-config` command in EXEC Privilege mode.

Monitoring RADIUS

To view information on RADIUS transactions, use the following command.

- View RADIUS transactions to troubleshoot problems.
EXEC Privilege mode

```
debug radius
```

TACACS+

The Dell Networking OS supports terminal access controller access control system (TACACS+) client, including support for login authentication.

Configuration Task List for TACACS+

The following list includes the configuration task for TACACS+ functions.

- [Choosing TACACS+ as the Authentication Method](#)
- [Monitoring TACACS+](#)
- [TACACS+ Remote Authentication and Authorization](#)
- [Specifying a TACACS+ Server Host](#)
- [Choosing TACACS+ as the Authentication Method](#)

For a complete listing of all commands related to TACACS+, refer to the *Security* chapter in the *Dell Networking OS Command Reference Guide*.

Choosing TACACS+ as the Authentication Method

One of the login authentication methods available is TACACS+ and the user's name and password are sent for authentication to the TACACS hosts specified.

To use TACACS+ to authenticate users, specify at least one TACACS+ server for the system to communicate with and configure TACACS+ as one of your authentication methods.

To select TACACS+ as the login authentication method, use the following commands.

1. Configure a TACACS+ server host.

CONFIGURATION mode

```
tacacs-server host {ip-address | host}
```

Enter the IP address or host name of the TACACS+ server.

Use this command multiple times to configure multiple TACACS+ server hosts.

2. Enter a text string (up to 16 characters long) as the name of the method list you wish to use with the TACACS+ authentication method.

CONFIGURATION mode

```
aaa authentication login {method-list-name | default} tacacs+ [...method3]
```

The TACACS+ method must not be the last method specified.

3. Enter LINE mode.

CONFIGURATION mode

```
line {aux 0 | console 0 | vty number [end-number]}
```

4. Assign the *method-list* to the terminal line.

LINE mode

```
login authentication {method-list-name | default}
```

Example of a Failed Authentication

To view the configuration, use the `show config` in LINE mode or the `show running-config tacacs +` command in EXEC Privilege mode.

If authentication fails using the primary method, the system employs the second method (or third method, if necessary) automatically. For example, if the TACACS+ server is reachable, but the server key is invalid, the system proceeds to the next authentication method. In the following example, the TACACS+ is incorrect, but the user is still authenticated by the secondary method.

First bold line: Server key purposely changed to incorrect value.

Second bold line: User authenticated using the secondary method.

```
Dell(conf)#
Dell(conf)#do show run aaa
!
aaa authentication enable default tacacs+ enable
aaa authentication enable LOCAL enable tacacs+
aaa authentication login default tacacs+ local
aaa authentication login LOCAL local tacacs+
aaa authorization exec default tacacs+ none
aaa authorization commands 1 default tacacs+ none
aaa authorization commands 15 default tacacs+ none
aaa accounting exec default start-stop tacacs+
aaa accounting commands 1 default start-stop tacacs+
aaa accounting commands 15 default start-stop tacacs+
Dell(conf)#
Dell(conf)#do show run tacacs+
!
tacacs-server key 7 d05206c308f4d35b
tacacs-server host 10.10.10.10 timeout 1
Dell(conf)#tacacs-server key angeline
Dell(conf)##RPM0-P:CP %SEC-5-LOGIN_SUCCESS: Login successful for user admin on
vty0 (10.11.9.209)
%RPM0-P:CP %SEC-3-AUTHENTICATION_ENABLE_SUCCESS: Enable password
authentication success on vty0 (10.11.9.209)
%RPM0-P:CP %SEC-5-LOGOUT: Exec session is terminated for user admin on line
vty0 (10.11.9.209)
Dell(conf)#username angeline password angeline
Dell(conf)##RPM0-P:CP %SEC-5-LOGIN_SUCCESS: Login successful for user angeline
```

```
on vty0 (10.11.9.209)
%RPM0-P:CP %SEC-3-AUTHENTICATION_ENABLE_SUCCESS: Enable password
authentication success on vty0 ( 10.11.9.209 )
```

Monitoring TACACS+

To view information on TACACS+ transactions, use the following command.

- View TACACS+ transactions to troubleshoot problems.

EXEC Privilege mode

```
debug tacacs+
```

TACACS+ Remote Authentication and Authorization

The Dell Networking OS takes the access class from the TACACS+ server. Access class is the class of service that restricts Telnet access and packet sizes.

If you have configured remote authorization, the system ignores the access class you have configured for the VTY line. The system instead gets this access class information from the TACACS+ server. The Dell Networking OS must know the username and password of the incoming user before it can fetch the access class from the server. A user, therefore, at least sees the login prompt. If the access class denies the connection, the system closes the Telnet session immediately.

The following example demonstrates how to configure the access-class from a TACACS+ server. This configuration ignores the configured access-class on the VTY line. If you have configured a deny10 ACL on the TACACS+ server, Dell downloads it and applies it. If the user is found to be coming from the 10.0.0.0 subnet, Dell also immediately closes the Telnet connection. Note, that no matter where the user is coming from, they see the login prompt.

When configuring a TACACS+ server host, you can set different communication parameters, such as the key password.

Example of Specifying a TACACS+ Server Host

```
Dell#
Dell(conf)#
Dell(conf)#ip access-list standard deny10
Dell(conf-std-nacl)#permit 10.0.0.0/8
Dell(conf-std-nacl)#deny any
Dell(conf)#
Dell(conf)#aaa authentication login tacacsmethod tacacs+
Dell(conf)#aaa authentication exec tacacsauthorization tacacs+
Dell(conf)#tacacs-server host 25.1.1.2 key Force10
Dell(conf)#
Dell(conf)#line vty 0 9
Dell(config-line-vty)#login authentication tacacsmethod
Dell(config-line-vty)#authorization exec tacacauthor
Dell(config-line-vty)#
Dell(config-line-vty)#access-class deny10
Dell(config-line-vty)#end
```

Specifying a TACACS+ Server Host

To specify a TACACS+ server host and configure its communication parameters, use the following command.

- Enter the host name or IP address of the TACACS+ server host.
CONFIGURATION mode

```
tacacs-server host {hostname | ip-address} [port port-number] [timeout seconds] [key key]
```

Configure the optional communication parameters for the specific host:

- `port port-number`: the range is from 0 to 65535. Enter a TCP port number. The default is **49**.
- `timeout seconds`: the range is from 0 to 1000. Default is **10 seconds**.
- `key key`: enter a string for the key. The key can be up to 42 characters long. This key must match a key configured on the TACACS+ server host. This parameter must be the last parameter you configure.

If you do not configure these optional parameters, the default global values are applied.

Example of Connecting with a TACACS+ Server Host

To specify multiple TACACS+ server hosts, configure the `tacacs-server host` command multiple times. If you configure multiple TACACS+ server hosts, the system attempts to connect with them in the order in which they were configured.

To view the TACACS+ configuration, use the `show running-config tacacs+` command in EXEC Privilege mode.

To delete a TACACS+ server host, use the `no tacacs-server host {hostname | ip-address}` command.

```
freebsd2# telnet 2200:2200:2200:2200:2200::2202
Trying 2200:2200:2200:2200:2200::2202...
Connected to 2200:2200:2200:2200:2200::2202.
Escape character is '^]'.
Login: admin
Password:
Dell#
```

Command Authorization

The AAA command authorization feature configures the Dell Networking OS to send each configuration command to a TACACS server for authorization before it is added to the running configuration.

By default, the AAA authorization commands configure the system to check both EXEC mode and CONFIGURATION mode commands. To enable only EXEC mode command checking, use the `no aaa authorization config-commands` command.

If rejected by the AAA server, the command is not added to the running config, and a message displays:

```
04:07:48: %RPM0-P:CP %SEC-3-SEC_AUTHORIZATION_FAIL: Authorization failure
Command
authorization failed for user (denyall) on vty0 ( 10.11.9.209 )
```

Protection from TCP Tiny and Overlapping Fragment Attacks

Tiny and overlapping fragment attack is a class of attack where configured ACL entries — denying TCP port-specific traffic — is bypassed and traffic is sent to its destination although denied by the ACL.


RFC 1858 and 3128 proposes a countermeasure to the problem. This countermeasure is configured into the line cards and enabled by default.

Enabling SCP and SSH

Secure shell (SSH) is a protocol for secure remote login and other secure network services over an insecure network. The Dell Networking OS is compatible with SSH versions 1.5 and 2, both the client and server modes. SSH sessions are encrypted and use authentication. Starting with Dell Networking OS Release 9.2(0.0), SSH is enabled by default.

For details about the command syntax, refer to the *Security* chapter in the *Dell Networking OS Command Line Interface Reference Guide*.

SCP is a remote file copy program that works with SSH and the Dell Networking OS supports.

 **NOTE:** The Windows-based WinSCP client software is not supported for secure copying between a PC and an Dell Networking OS-based system. Unix-based SCP client software is supported.

To use the SSH client, use the following command.

- Open an SSH connection and specifying the host name, username, port number, encryption cipher, HMAC algorithm, and version of the SSH client.

EXEC Privilege mode

```
ssh {hostname} [-l username | -p port-number | -v {1 | 2} | -c encryption
cipher | -m HMAC algorithm
```

hostname is the IP address or host name of the remote device. Enter an IPv4 or IPv6 address in dotted decimal format (A.B.C.D).

- Configure the Dell Networking system as an SCP/SSH server.

CONFIGURATION mode

```
ip ssh server {enable | port port-number | version | vrf}
```

- Configure the Dell Networking system as an SSH server that uses only version 1 or 2.

CONFIGURATION mode

```
ip ssh server version {1|2}
```

- Display SSH connection information.

EXEC Privilege mode

```
show ip ssh
```

Specifying an SSH Version

The following example shows using the `ip ssh server version 2` command to enable SSH version 2 and the `show ip ssh` command to confirm the setting.

```
Dell(conf)#ip ssh server version 2
Dell(conf)#do show ip ssh
SSH server          : disabled.
SSH server version  : v2.
Password Authentication : enabled.
Hostbased Authentication : disabled.
RSA Authentication   : disabled.
```

To disable SSH server functions, use the `no ip ssh server enable` command.

Using SCP with SSH to Copy a Software Image

To use secure copy (SCP) to copy a software image through an SSH connection from one switch to another, use the following commands.

1. On Chassis One, set the SSH port number (**port 22** by default).
CONFIGURATION mode

```
ip ssh server port number
```
2. On Chassis One, enable SSH.
CONFIGURATION mode

```
ip ssh server enable
```
3. On Chassis Two, invoke SCP.
CONFIGURATION mode

```
copy scp: flash:
```
4. On Chassis Two, in response to prompts, enter the path to the desired file and enter the port number specified in Step 1.
EXEC Privilege mode

Example of Using SCP to Copy from an SSH Server on Another Switch

Other SSH-related commands include:

- `crypto key generate`: generate keys for the SSH server.
- `debug ip ssh`: enables collecting SSH debug information.
- `ip scp topdir`: identify a location for files used in secure copy transfer.
- `ip ssh authentication-retries`: configure the maximum number of attempts that should be used to authenticate a user.
- `ip ssh connection-rate-limit`: configure the maximum number of incoming SSH connections per minute.
- `ip ssh hostbased-authentication enable`: enable host-based authentication for the SSHv2 server.
- `ip ssh key-size`: configure the size of the server-generated RSA SSHv1 key.

- `ip ssh password-authentication enable`: enable password authentication for the SSH server.
- `ip ssh pub-key-file`: specify the file the host-based authentication uses.
- `ip ssh rhostsfile`: specify the rhost file the host-based authorization uses.
- `ip ssh rsa-authentication enable`: enable RSA authentication for the SSHv2 server.
- `ip ssh rsa-authentication`: add keys for the RSA authentication.
- `show crypto`: display the public part of the SSH host-keys.
- `show ip ssh client-pub-keys`: display the client public keys used in host-based authentication.
- `show ip ssh rsa-authentication`: display the authorized-keys for the RSA authentication.

The following example shows the use of SCP and SSH to copy a software image from one switch running SSH server on UDP port 99 to the local switch.

```
Dell#copy scp: flash:
Address or name of remote host []: 10.10.10.1
Port number of the server [22]: 99
Source file name []: test.cfg
User name to login remote host: admin
Password to login remote host:
```

Removing the RSA Host Keys and Zeroizing Storage

Use the `crypto key zeroize rsa` command to delete the host key pairs, both the public and private key information for RSA 1 and or RSA 2 types. Note that when FIPS mode is enabled there is no RSA 1 key pair. Any memory currently holding these keys is zeroized (written over with zeroes) and the NVRAM location where the keys are stored for persistence across reboots is also zeroized.

To remove the generated RSA host keys and zeroize the key storage location, use the `crypto key zeroize rsa` command in CONFIGURATION mode.

```
Dell(conf)#crypto key zeroize rsa
```

Configuring When to Re-generate an SSH Key

You can configure the time-based or volume-based rekey threshold for an SSH session. If both threshold types are configured, the session rekeys when either one of the thresholds is reached.

To configure the time or volume rekey threshold at which to re-generate the SSH key during an SSH session, use the `ip ssh rekey [time rekey-interval] [volume rekey-limit]` command. CONFIGURATION mode.

Configure the following parameters:

- *rekey-interval*: time-based rekey threshold for an SSH session. The range is from 10 to 1440 minutes. The default is **60** minutes.
- *rekey-limit*: volume-based rekey threshold for an SSH session. The range is from 1 to 4096 to megabytes. The default is **1024** megabytes.

Examples

The following example configures the time-based rekey threshold for an SSH session to 30 minutes.

```
Dell(conf)#ip ssh rekey time 30
```

The following example configures the volume-based rekey threshold for an SSH session to 4096 megabytes.

```
Dell(conf)#ip ssh rekey volume 4096
```

Configuring the SSH Server Key Exchange Algorithm

To configure the key exchange algorithm for the SSH server, use the `ip ssh server kex key-exchange-algorithm` command in CONFIGURATION mode.

key-exchange-algorithm : Enter a space-delimited list of key exchange algorithms that will be used by the SSH server.

The following key exchange algorithms are available:

- diffie-hellman-group-exchange-sha1
- diffie-hellman-group1-sha1
- diffie-hellman-group14-sha1

The default key exchange algorithms are the following:

- diffie-hellman-group-exchange-sha1
- diffie-hellman-group1-sha1
- diffie-hellman-group14-sha1

When FIPS is enabled, the default is diffie-hellman-group14-sha1.

Example of Configuring a Key Exchange Algorithm

The following example shows you how to configure a key exchange algorithm.

```
Dell(conf)# ip ssh server kex diffie-hellman-group-exchange-sha1 diffie-hellman-group14-sha1
```

Configuring the HMAC Algorithm for the SSH Server

To configure the HMAC algorithm for the SSH server, use the `ip ssh server mac hmac-algorithm` command in CONFIGURATION mode.

hmac-algorithm: Enter a space-delimited list of keyed-hash message authentication code (HMAC) algorithms supported by the SSH server.

The following HMAC algorithms are available:

- hmac-md5
- hmac-md5-96
- hmac-sha1

- hmac-sha1-96
- hmac-sha2-256
- hmac-sha2-256-96

The default HMAC algorithms are the following:

- hmac-md5
- hmac-md5-96
- hmac-sha1
- hmac-sha1-96
- hmac-sha2-256
- hmac-sha2-256-96

When FIPS is enabled, the default HMAC algorithm is hmac-sha1-96.

Example of Configuring a HMAC Algorithm

The following example shows you how to configure a HMAC algorithm list.

```
Dell(conf)# ip ssh server mac hmac-sha1-96
```

Configuring the SSH Server Cipher List

To configure the cipher list supported by the SSH server, use the `ip ssh server cipher cipher-list` command in CONFIGURATION mode.

cipher-list:- Enter a space-delimited list of ciphers the SSH server will support.

The following ciphers are available.

- 3des-cbc
- aes128-cbc
- aes192-cbc
- aes256-cbc
- aes128-ctr
- aes192-ctr
- aes256-ctr

The default cipher list is 3des-cbc,aes128-cbc,aes192-cbc,aes256-cbc,aes128-ctr,aes192-ctr,aes256-ctr

Example of Configuring a Cipher List

The following example shows you how to configure a cipher list.

```
Dell(conf)#ip ssh server cipher 3des-cbc aes128-cbc aes128-ctr
```


Secure Shell Authentication

Secure Shell (SSH) is disabled by default.

Enable SSH using the `ip ssh server enable` command.

SSH supports three methods of authentication:

- [Enabling SSH Authentication by Password](#)
- [Using RSA Authentication of SSH](#)
- [Configuring Host-Based SSH Authentication](#)

Important Points to Remember

- If you enable more than one method, the order in which the methods are preferred is based on the `ssh_config` file on the Unix machine.
- When you enable all the three authentication methods, password authentication is the backup method when the RSA method fails.
- The files `known_hosts` and `known_hosts2` are generated when a user tries to SSH using version 1 or version 2, respectively.

Enabling SSH Authentication by Password

Authenticate an SSH client by prompting for a password when attempting to connect to the Dell Networking system. This setup is the simplest method of authentication and uses SSH version 1. To enable SSH password authentication, use the following command.

- Enable SSH password authentication.
CONFIGURATION mode

```
ip ssh password-authentication enable
```

Example of Enabling SSH Password Authentication

To view your SSH configuration, use the `show ip ssh` command from EXEC Privilege mode.

```
Dell(conf)#ip ssh server enable
% Please wait while SSH Daemon initializes ... done.
Dell(conf)#ip ssh password-authentication enable
Dell#sh ip ssh
SSH server           : enabled.
Password Authentication : enabled.
Hostbased Authentication : disabled.
RSA Authentication    : disabled.
Vty Encryption
```

Using RSA Authentication of SSH

The following procedure authenticates an SSH client based on an RSA key using RSA authentication. This method uses SSH version 2.

1. On the SSH client (Unix machine), generate an RSA key, as shown in the following example.
2. Copy the public key `id_rsa.pub` to the Dell Networking system.
3. Disable password authentication if enabled.

CONFIGURATION mode

```
no ip ssh password-authentication enable
```

4. Bind the public keys to RSA authentication.

EXEC Privilege mode

```
ip ssh rsa-authentication enable
```

5. Bind the public keys to RSA authentication.

EXEC Privilege mode

```
ip ssh rsa-authentication my-authorized-keys flash://public_key
```

Example of Generating RSA Keys

```
admin@Unix_client#ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/home/admin/.ssh/id_rsa):
/home/admin/.ssh/id_rsa already exists.
Overwrite (y/n)? y
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/admin/.ssh/id_rsa.
Your public key has been saved in /home/admin/.ssh/id_rsa.pub.
```

Configuring Host-Based SSH Authentication

Authenticate a particular host. This method uses SSH version 2.

To configure host-based authentication, use the following commands.

1. Configure RSA Authentication. Refer to [Using RSA Authentication of SSH](#).
2. Create *shosts* by copying the public RSA key to the file *shosts* in the directory *.ssh*, and write the IP address of the host to the file.

```
cp /etc/ssh/ssh_host_rsa_key.pub /.ssh/shosts
```

Refer to the first example.

3. Create a list of IP addresses and usernames that are permitted to SSH in a file called *rhosts*. Refer to the second example.

4. Copy the file *shosts* and *rhosts* to the Dell Networking system.
5. Disable password authentication and RSA authentication, if configured

CONFIGURATION mode or EXEC Privilege mode

```
no ip ssh password-authentication Or no ip ssh rsa-authentication
```

6. Enable host-based authentication.

CONFIGURATION mode

```
ip ssh hostbased-authentication enable
```

7. Bind *shosts* and *rhosts* to host-based authentication.

CONFIGURATION mode

```
ip ssh pub-key-file flash://filename Or ip ssh rhostsfile flash://filename
```

Example of Creating *shosts*

Example of Creating *rhosts*

```
admin@Unix_client# cd /etc/ssh

admin@Unix_client# ls
moduli      sshd_config  ssh_host_dsa_key.pub  ssh_host_key.pub
ssh_host_rsa_key.pub  ssh_config  ssh_host_dsa_key  ssh_host_key
ssh_host_rsa_key

admin@Unix_client# cat ssh_host_rsa_key.pub

ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEA8K7jLZRVfjgHJzUOmXxuIbZx/
AyWhVgJDQh39k8v3e8eQvLnHBIsqIL8jVy1QHhUeb7GaDlJVEDAMz30myqQbJgXBBRTWgBpLWwL/
doyUXFufjiL9YmoVTkbKcFmxJEMkE3JyHanEi7hg34LChjk9hL1by8cYZP2kYS2lnSyQWk=

admin@Unix_client# ls

id_rsa id_rsa.pub shosts

admin@Unix_client# cat shosts

10.16.127.201, ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAIEA8K7jLZRVfjgHJzUOmXxuIbZx/AyW
hVgJDQh39k8v3e8eQvLnHBIsqIL8jVy1QHhUeb7GaDlJVEDAMz30myqQbJgXBBRTWgBpLWwL/

admin@Unix_client# ls
id_rsa id_rsa.pub rhosts shosts
admin@Unix_client# cat rhosts
10.16.127.201 admin
```

Using Client-Based SSH Authentication

To SSH from the chassis to the SSH client, use the following command.

This method uses SSH version 1 or version 2. If the SSH port is a non-default value, use the `ip ssh server port number` command to change the default port number. You may only change the port number when SSH is disabled. Then use the `-p` option with the `ssh` command.

- SSH from the chassis to the SSH client.

```
ssh ip_address
```

Example of Client-Based SSH Authentication

```
Dell#ssh 10.16.127.201 ?
-c Encryption cipher to use (for v2 clients only)
-l User name option
-m HMAC algorithm to use (for v2 clients only)
-p SSH server port option (default 22)
-v SSH protocol version
```

Troubleshooting SSH

To troubleshoot SSH, use the following information.

You may not bind `id_rsa.pub` to RSA authentication while logged in via the console. In this case, this message displays:`%Error: No username set for this term.`

Enable host-based authentication on the server (Dell Networking system) and the client (Unix machine). The following message appears if you attempt to log in via SSH and host-based is disabled on the client.

In this case, verify that host-based authentication is set to “Yes” in the *file ssh_config* (root permission is required to edit this file): `permission denied (host based)`.

If the IP address in the RSA key does not match the IP address from which you attempt to log in, the following message appears. In this case, verify that the name and IP address of the client is contained in the *file /etc/hosts*: `RSA Authentication Error`.

Telnet

To use Telnet with SSH, first enable SSH, as previously described.

By default, the Telnet daemon is enabled. If you want to disable the Telnet daemon, use the following command, or disable Telnet in the startup config. To enable or disable the Telnet daemon, use the `[no] ip telnet server enable` command.

Example of Using Telnet for Remote Login

```
Dell(conf)#ip telnet server enable
Dell(conf)#no ip telnet server enable
```

VTY Line and Access-Class Configuration

Various methods are available to restrict VTY access in the Dell Networking OS. These depend on which authentication scheme you use — line, local, or remote.

Table 56. VTY Access

Authentication Method	VTY access-class support?	Username access-class support?	Remote authorization support?
Line	YES	NO	NO
Local	NO	YES	NO
TACACS+	YES	NO	YES (with the Dell Networking OS version 5.2.1.0 and later)
RADIUS	YES	NO	YES (with the Dell Networking OS version 6.1.1.0 and later)

The Dell Networking OS provides several ways to configure access classes for VTY lines, including:

- [VTY Line Local Authentication and Authorization](#)
- [VTY Line Remote Authentication and Authorization](#)

VTY Line Local Authentication and Authorization

Dell Networking OS retrieves the access class from the local database.


To use this feature:

1. Create a username.
2. Enter a password.
3. Assign an access class.
4. Enter a privilege level.


You can assign line authentication on a per-VTY basis; it is a simple password authentication, using an access-class as authorization.

Configure local authentication globally and configure access classes on a per-user basis.

The Dell Networking OS can assign different access classes to different users by username. Until users attempt to log in, the system does not know if they will be assigned a VTY line. This means that incoming users always see a login prompt even if you have excluded them from the VTY line with a deny-all access class. After users identify themselves, the system retrieves the access class from the local database and applies it. (The Dell Networking OS then can close the connection if a user is denied access.)

 **NOTE:** If a VTY user logs in with RADIUS authentication, the privilege level is applied from the RADIUS server only if you configure RADIUS authentication.

The following example shows how to allow or deny a Telnet connection to a user. Users see a login prompt even if they cannot log in. No access class is configured for the VTY line. It defaults from the local database.

 **NOTE:** For more information, refer to [Access Control Lists \(ACLs\)](#).

Example of Configuring VTY Authorization Based on Access Class Retrieved from a Local Database (Per User)

```
Dell(conf)#user gooduser password abc privilege 10 access-class permitall
Dell(conf)#user baduser password abc privilege 10 access-class denyall
Dell(conf)#
Dell(conf)#aaa authentication login localmethod local
Dell(conf)#
Dell(conf)#line vty 0 9
Dell(config-line-vty)#login authentication localmethod
Dell(config-line-vty)#end
```

VTY Line Remote Authentication and Authorization

The Dell Networking OS retrieves the access class from the VTY line.

The Dell Networking OS takes the access class from the VTY line and applies it to ALL users. The system does not need to know the identity of the incoming user and can immediately apply the access class. If the authentication method is RADIUS, TACACS+, or line, and you have configured an access class for the VTY line, the system immediately applies it. If the access-class is set to deny all or deny for the incoming subnet, the system closes the connection without displaying the login prompt. The following example shows how to deny incoming connections from subnet 10.0.0.0 without displaying a login prompt. The example uses TACACS+ as the authentication mechanism.

Example of Configuring VTY Authorization Based on Access Class Retrieved from the Line (Per Network Address)

```
Dell(conf)#ip access-list standard deny10
Dell(conf-ext-nacl)#permit 10.0.0.0/8
Dell(conf-ext-nacl)#deny any
```

```
Dell(conf)#
Dell(conf)#aaa authentication login tacacsmethod tacacs+
Dell(conf)#tacacs-server host 256.1.1.2 key Force10
Dell(conf)#
Dell(conf)#line vty 0 9
Dell(config-line-vty)#login authentication tacacsmethod
Dell(config-line-vty)#
Dell(config-line-vty)#access-class deny10
Dell(config-line-vty)#end
(same applies for radius and line authentication)
```

VTY MAC-SA Filter Support

The Dell Networking OS supports MAC access lists which permit or deny users based on their source MAC address.

With this approach, you can implement a security policy based on the source MAC address.

To apply a MAC ACL on a VTY line, use the same `access-class` command as IP ACLs.

The following example shows how to deny incoming connections from subnet 10.0.0.0 without displaying a login prompt.

Example of Configuring VTY Authorization Based on MAC ACL for the Line (Per MAC Address)

```
Dell(conf)#mac access-list standard sourcemac
Dell(config-std-mac)#permit 00:00:5e:00:01:01
Dell(config-std-mac)#deny any
Dell(conf)#
Dell(conf)#line vty 0 9
Dell(config-line-vty)#access-class sourcemac
Dell(config-line-vty)#end
```

Role-Based Access Control

With Role-Based Access Control (RBAC), access and authorization is controlled based on a user's role. Users are granted permissions based on their user roles, not on their individual user ID. User roles are created for job functions and through those roles they acquire the permissions to perform their associated job function.

This chapter consists of the following sections:

- Overview
- Privilege-or-Role Mode Versus Role-only Mode
- Configuring Role-based Only AAA Authorization
- System-Defined RBAC User Roles
- Creating a New User Role
- Modifying Command Permissions for Roles
- Adding and Deleting Users from a Role
- Role Accounting
- Configuring AAA Authentication for Roles

- Configuring AAA Authorization for Roles
- Configuring an Accounting for Roles
- Applying an Accounting Method to a Role
- Displaying Active Accounting Sessions for Roles
- Configuring TACACS+ and RADIUS VSA Attributes for RBAC
- Displaying User Roles
- Displaying Accounting for User Roles
- Displaying Information About Roles Logged into the Switch
- Display Role Permissions Assigned to a Command


Overview of RBAC

With Role-Based Access Control (RBAC), access and authorization is controlled based on a user's role. Users are granted permissions based on their user roles, not on their individual user ID. User roles are created for job functions and through those roles they acquire the permissions to perform their associated job function. Each user can be assigned only a single role. Many users can have the same role.

The Dell Networking OS supports the constrained RBAC model. With a constrained RBAC model, you can inherit permissions when you create a new user role, restrict or add commands a user can enter and the actions the user can perform. This allows for greater flexibility in assigning permissions for each command to each role and as a result, it is easier and much more efficient to administer user rights. If a user's role matches one of the allowed user roles for that command, then command authorization is granted.

A constrained RBAC model provides for separation of duty and as a result, provides greater security than the hierarchical RBAC model. Essentially, a constrained model puts some limitations around each role's permissions to allow you to partition of tasks. However, some inheritance is possible.

Default command permissions are based on CLI mode (such as configure, interface, router), any specific command settings, and the permissions allowed by the privilege and role commands. The role command allows you to change permissions based on the role. You can modify the permissions specific to that command and/or command option. For more information, see [Modifying Command Permissions for Roles](#).

 **NOTE:** When you enter a user role, you have already been authenticated and authorized. You do not need to enter an enable password because you will be automatically placed in EXEC Priv mode.

For greater security, the ability to view event, audit, and security system log is associated with user roles. For information about these topics, see [Audit and Security Logs](#).

Privilege-or-Role Mode versus Role-only Mode

By default, the system provides access to commands determined by the user's role or by the user's privilege level. The user's role takes precedence over a user's privilege level. If the system is in "privilege or role" mode, then all existing user IDs can continue to access the switch even if they do not have a user role defined. To change to more secure mode, use role-based AAA authorization. When role-based only AAA authorization is configured, access to commands is determined only by the user's role. For more information, see [Configuring Role-based Only AAA Authorization](#).

Configuring Role-based Only AAA Authorization

You can configure authorization so that access to commands is determined only by the user's role. If the user has no user role, access to the system is denied as the user will not be able to login successfully. When you enable role-based only AAA authorization using the **aaa authorization role-only** command in Configuration mode, the Dell Networking OS checks to ensure that you do not lock yourself out and that the user authentication is available for all terminal lines.

Pre-requisites

Before you enable role-based only AAA authorization:


1. Locally define a system administrator user role. This will give you access to login with full permissions even if network connectivity to remote authentication servers is not available.
2. Configure login authentication on the console. This ensures that all users are properly identified through authentication no matter the access point.

If you do not configure login the authentication on the console, the system displays an error when you attempt to enable role-based only AAA authorization.

3. Specify an authentication method list (RADIUS, TACACS+, or Local).

You must specify at least local authentication. For consistency, the best practice is to define the same authentication method list across all lines, in the same order of comparison; for example VTY and console port.

You could also use the default authentication method to apply to all the LINES (console port, VTY).

 **NOTE:** The authentication method list should be in the same order as the authorization method list. For example, if you configure the authentication method list in the following order (TACACS+, local), Dell Networking recommends that authorization method list is configured in the same order (TACACS+, local).

4. Specify authorization method list (RADIUS, TACACS+, or Local). You must at least specify local authorization.

For consistency, the best practice is to define the same authorization method list across all lines, in the same order of comparison; for example VTY and console port.

You could also use the default authorization method list to apply to all the LINES (console port, VTY).

If you do not, the following error is displayed when you attempt to enable role-based only AAA authorization.

```
% Error: Exec authorization must be applied to more than one line to be useful, e.g. console and vty lines. Could use default authorization method list as alternative.
```

5. Verify the configuration has been applied to the console or VTY line.

```
Dell (conf)#do show running-config line
!
line console 0
login authentication test
authorization exec test
```



```
exec-timeout 0 0
line vty 0
login authentication test
authorization exec test
line vty 1
login authentication test
authorization exec test
```

To enable role-based only AAA authorization:

```
Dell(conf)#aaa authorization role-only
```

System-Defined RBAC User Roles

By default, the Dell Networking OS provides 4 system defined user roles. You can create up to 8 additional user roles.

 **NOTE:** You cannot delete any system defined roles.

The system defined user roles are as follows:

- Network Operator (netoperator) - This user role has no privilege to modify any configuration on the switch. You can access Exec mode (monitoring) to view the current configuration and status information.
- Network Administrator (netadmin): This user role can configure, display, and debug the network operations on the switch. You can access all of the commands that are available from the network operator user role. This role does not have access to the commands that are available to the system security administrator for cryptography operations, AAA, or the commands reserved solely for the system administrator.
- Security Administrator (secadmin): This user role can control the security policy across the systems that are within a domain or network topology. The security administrator commands include FIPS mode enablement, password policies, inactivity timeouts, banner establishment, and cryptographic key operations for secure access paths.
- System Administrator (sysadmin). This role has full access to all the commands in the system, exclusive access to commands that manipulate the file system formatting, and access to the system shell. This role can also create user IDs and user roles.

The following summarizes the modes that the predefined user roles can access.

Role	Modes
netoperator	
netadmin	Exec Config Interface Router IP Route-map Protocol MAC
secadmin	Exec Config Line
sysadmin	Exec Config Interface Line Router IP Route-map Protocol MAC


User Roles

This section describes how to create a new user role and configure command permissions and contains the following topics.

- [Creating a New User Role](#)
- [Modifying Command Permissions for Roles](#)
- [Adding and Deleting Users from a Role](#)

Creating a New User Role

Instead of using the system defined user roles, you can create a new user role that best matches your organization. When you create a new user role, you can first inherit permissions from one of the system defined roles. Otherwise you would have to create a user role's command permissions from scratch. You then restrict commands or add commands to that role. For more information about this topic, see [Modifying Command Permissions for Roles](#).

 **NOTE:** You can change user role permissions on system pre-defined user roles or user-defined user roles.

Important Points to Remember

Consider the following when creating a user role:

- Only the system administrator and user-defined roles inherited from the system administrator can create roles and user names. Only the system administrator, security administrator, and roles inherited from these can use the "role" command to modify command permissions. The security administrator and roles inherited by security administrator can only modify permissions for commands they already have access to.
- Make sure you select the correct role you want to inherit.
- If you inherit a user role, you cannot modify or delete the inheritance. If you want to change or remove the inheritance, delete the user role and create it again. If the user role is in use, you cannot delete the user role.

1. Create a new user role
CONFIGURATION mode

```
userrole name [inherit existing-role-name]
```

2. Verify that the new user role has inherited the security administrator permissions.

```
Dell(conf)#do show userroles
```

```
EXEC Privilege mode
```

3. After you create a user role, configure permissions for the new user role. See [Modifying Command Permissions for Roles](#).

Example of Creating a User Role

The configuration in the following example creates a new user role, **myrole**, which inherits the security administrator (secadmin) permissions.

Create a new user role, **myrole** and inherit security administrator permissions.

```
Dell(conf)#userrole myrole inherit secadmin
```

Verify that the user role, **myrole**, has inherited the security administrator permissions. The output highlighted in **bold** indicates that the user role has successfully inherited the security administrator permissions.

```
Dell(conf)#do show userroles
```


```
***** Mon Apr 28 14:46:25 PDT 2014 *****

Authorization Mode:  role or privilege
Role      Inheritance  Modes
netoperator

netadmin          Exec Config Interface Router IP Route-map Protocol MAC
secadmin          Exec Config Line
sysadmin          Exec Config Interface Line Router IP Route-map
Protocol MAC.
myrole            secadmin      Exec Config Line
```

Modifying Command Permissions for Roles

You can modify (add or delete) command permissions for newly created user roles and system defined roles using the `role mode { { { addrole | deleterole } role-name } | reset } command` command in Configuration mode.

 **NOTE:** You cannot modify system administrator command permissions.

If you add or delete command permissions using the `role` command, those changes only apply to the specific user role. They do not apply to other roles that have inheritance from that role. Authorization and accounting only apply to the roles specified in that configuration.

When you modify a command for a role, you specify the role, the mode, and whether you want to restrict access using the `deleterole` keyword or grant access using the `addrole` keyword followed by the command you are controlling access. For information about how to create new roles, see also [Creating a New User Role](#).

The following output displays the modes available for the `role` command.

```
Dell (conf)#role ?
configure      Global configuration mode
exec           Exec Mode
interface      Interface configuration mode
line           Line Configuration mode
route-map      Route map configuration mode
router         Router configuration mode
```

Examples: Deny Network Administrator from Using the show users Command.

The following example denies the `netadmin` role from using the `show users` command and then verifies that `netadmin` cannot access the `show users` command in `exec` mode. Note that the

netadmin role is not listed in the Role access: secadmin, sysadmin, which means the netadmin cannot access the show users command.

```
Dell(conf)#role exec deleterole netadmin show users
```

```
Dell#show role mode exec show users
Role access: secadmin, sysadmin
```

Example: Allow Security Administrator to Configure Spanning Tree

The following example allows the security administrator (secadmin) to configure the spanning tree protocol. Note *command* is protocol spanning-tree.

```
Dell(conf)#role configure addrole secadmin protocol spanning-tree
```

Example: Allow Security Administrator to Access Interface Mode

The following example allows the security administrator (secadmin) to access Interface mode.

```
Dell(conf)#role configure addrole secadmin ?
LINE      Initial keywords of the command to modify
Dell(conf)#role configure addrole secadmin interface
```

Example: Allow Security Administrator to Access Only 10-Gigabit Ethernet Interfaces

The following example allows the security administrator (secadmin) to only access 10-Gigabit Ethernet interfaces and then shows that the secadmin, highlighted in bold, can now access Interface mode. However, the secadmin can only access 10-Gigabit Ethernet interfaces.

```
Dell(conf)#role configure addrole secadmin ?
LINE      Initial keywords of the command to modify
Dell(conf)#role configure addrole secadmin interface tengigabitethernet
```

```
Dell(conf)#show role mode configure interface
Role access: netadmin, secadmin, sysadmin
```

Example: Verify that the Security Administrator Can Access Interface Mode

The following example shows that the secadmin role can now access Interface mode (highlighted in bold).

Role	Inheritance	Modes
netoperator		
netadmin		Exec Config Interface Router IP RouteMap Protocol MAC
secadmin		Exec Config Interface
Line		
sysadmin		Exec Config Interface Line Router IP RouteMap Protocol MAC

Example: Remove Security Administrator Access to Line Mode.

The following example removes the secadmin access to LINE mode and then verifies that the security administrator can no longer access LINE mode, using the show role mode configure line command in EXEC Privilege mode.

```
Dell(conf)#role configure deleterole secadmin ?
LINE      Initial keywords of the command to modify
Dell(conf)#role configure deleterole secadmin line
```

```
Dell(conf)#do show role mode ?
configure                               Global configuration mode
exec                                     Exec Mode
interface                               Interface configuration mode
line                                    Line Configuration mode
route-map                               Route map configuration mode
router                                  Router configuration mode
```

```
Dell(conf)#do show role mode configure line
Role access:sysadmin
```

Example: Grant and Remove Security Administrator Access to Configure Protocols

By default, the system defined role, `secadmin`, is not allowed to configure protocols. The following example first grants the `secadmin` role to configure protocols and then removes access to configure protocols.

```
Dell(conf)#role configure addrole secadmin protocol
Dell(conf)#role configure deleterole secadmin protocol
```

Example: Resets Only the Security Administrator role to its original setting.

The following example resets only the `secadmin` role to its original setting.

```
Dell(conf)#no role configure addrole secadmin protocol
```

Example: Reset System-Defined Roles and Roles that Inherit Permissions

In the following example the command protocol permissions are reset to their original setting or one or more of the system-defined roles and any roles that inherited permissions from them.

```
Dell(conf)#role configure reset protocol
```

Adding and Deleting Users from a Role


To create a user name that is authenticated based on a user role, use the `username name password encryption-type password role role-name` command in CONFIGURATION mode.

Example

The following example creates a user name that is authenticated based on a user role.

```
Dell (conf) #username john password 0 password role secadmin
```

The following example deletes a user role.

 **NOTE:** If you already have a user ID that exists with a privilege level, you can add the user role to username that has a privilege

```
Dell (conf) #no username john
```

The following example adds a user, to the `secadmin` user role.

```
Dell (conf)#username john role secadmin password 0 password
```

AAA Authentication and Authorization for Roles

This section describes how to configure AAA Authentication and Authorization for Roles.

Configuration Task List for AAA Authentication and Authorization for Roles


This section contains the following AAA Authentication and Authorization for Roles configuration tasks:

- [Configuring AAA Authentication for Roles](#)
- [Configuring AAA Authorization for Roles](#)
- [Configuring TACACS+ and RADIUS VSA Attributes for RBAC](#)

Configure AAA Authentication for Roles

Authentication services verify the user ID and password combination. Users with defined roles and users with privileges are authenticated with the same mechanism. There are six methods available for authentication: **radius**, **tacacs+**, **local**, **enable**, **line**, and **none**.

When role-based only AAA authorization is enabled, the **enable**, **line**, and **none** methods are not available. Each of these three methods allows users to be verified with either a password that is not specific to their user ID or with no password at all. Because of the lack of security these methods are not available for role only mode. When the system is in role-only mode, users that have only privilege levels are denied access to the system because they do not have a role. For information about role only mode, see [Configuring Role-based Only AAA Authorization](#).

 **NOTE:** Authentication services only validate the user ID and password combination. To determine which commands are permitted for users, configure authorization. For information about how to configure authorization for roles, see [Configure AAA Authorization for Roles](#).

To configure AAA authentication, use the **aaa authentication** command in CONFIGURATION mode.

```
aaa authentication login {method-list-name | default} method [... method4]
```

Configure AAA Authorization for Roles

Authorization services determine if the user has permission to use a command in the CLI. Users with only privilege levels can use commands in privilege-or-role mode (the default) provided their privilege level is the same or greater than the privilege level of those commands. Users with defined roles can use commands provided their role is permitted to use those commands. Role inheritance is also used to determine authorization.

Users with roles and privileges are authorized with the same mechanism. There are six methods available for authorization: **radius**, **tacacs+**, **local**, **enable**, **line**, and **none**.

When role-based only AAA authorization is enabled, the **enable**, **line**, and **none** methods are not available. Each of these three methods allows users to be authorized with either a password that is not specific to their userid or with no password at all. Because of the lack of security, these methods are not available for role-based only mode.

To configure AAA authorization, use the `aaa authorization exec` command in CONFIGURATION mode. The `aaa authorization exec` command determines which CLI mode the user will start in for their session; for example, Exec mode or Exec Privilege mode. For information about how to configure authentication for roles, see [Configure AAA Authentication for Roles](#).

```
aaa authorization exec {method-list-name | default} method [... method4]
```

You can further restrict users' permissions, using the `aaa authorization command` command in CONFIGURATION mode.


```
aaa authorization command {method-list-name | default} method [... method4]
```

Examples of Applying a Method List

The following configuration example applies a method list: TACACS+, RADIUS and local:

```
!  
radius-server host 10.16.150.203 key <clear-text>  
!  
tacacs-server host 10.16.150.203 key <clear-text>  
!  
aaa authentication login ucraaa tacacs+ radius local  
aaa authorization exec ucraaa tacacs+ radius local  
aaa accounting commands role netadmin ucraaa start-stop tacacs+  
!
```

The following configuration example applies a method list other than default to each VTY line.

 **NOTE:** Note that the methods were not applied to the console so the default methods (if configured) are applied there.

```
!  
line console 0  
exec-timeout 0 0  
line vty 0  
login authentication ucraaa  
authorization exec ucraaa  
accounting commands role netadmin ucraaa  
line vty 1  
login authentication ucraaa  
authorization exec ucraaa  
accounting commands role netadmin ucraaa  
line vty 2  
login authentication ucraaa  
authorization exec ucraaa  
accounting commands role netadmin ucraaa  
line vty 3  
login authentication ucraaa  
authorization exec ucraaa  
accounting commands role netadmin ucraaa  
line vty 4  
login authentication ucraaa  
authorization exec ucraaa  
accounting commands role netadmin ucraaa  
line vty 5  
login authentication ucraaa  
authorization exec ucraaa  
accounting commands role netadmin ucraaa  
line vty 6  
login authentication ucraaa  
authorization exec ucraaa  
accounting commands role netadmin ucraaa  
line vty 7  
login authentication ucraaa  
authorization exec ucraaa  
accounting commands role netadmin ucraaa  
line vty 8  
login authentication ucraaa
```

```

authorization exec ucraaa
accounting commands role netadmin ucraaa
line vty 9
login authentication ucraaa
authorization exec ucraaa
accounting commands role netadmin ucraaa
!
```

Configuring TACACS+ and RADIUS VSA Attributes for RBAC

For RBAC and privilege levels, the Dell Networking OS RADIUS and TACACS+ implementation supports two vendor-specific options: privilege level and roles. The Dell Networking vendor-ID is 6027 and the supported option has attribute of type string, which is titled "Force10-avpair". The value is a string in the following format:

```
protocol : attribute sep value
```

"attribute" and "value" are an attribute-value (AV) pair defined in the Dell Network OS TACACS+ specification, and "sep" is "=". These attributes allow the full set of features available for TACACS+ authorization and are authorized with the same attributes for RADIUS.

Example for Configuring a VSA Attribute for a Privilege Level 15

The following example configures an AV pair which allows a user to login from a network access server with a privilege level of 15, to have access to EXEC commands.

The format to create a Dell Network OS AV pair for privilege level is `shell:priv-lvl=<number>` where number is a value between 0 and 15.

```
Force10-avpair= "shell:priv-lvl=15"
```

Example for Creating a AVP Pair for System Defined or User-Defined Role

The following section shows you how to create an AV pair to allow a user to login from a network access server to have access to commands based on the user's role. The format to create an AV pair for a user role is `Force10-avpair= "shell:role=<user-role>"` where *user-role* is a user defined or system-defined role.

In the following example, you create an AV pair for a system-defined role, sysadmin.

```
Force10-avpair= "shell:role=sysadmin"
```

In the following example, you create an AV pair for a user-defined role. You must also define a role, using the `userrole myrole inherit` command on the switch to associate it with this AV pair.

```
Force10-avpair= "shell:role=myrole"
```

The string, "myrole", is associated with a TACACS+ user group. The user IDs are associated with the user group.

Role Accounting

This section describes how to configure role accounting and how to display active sessions for roles.

This sections consists of the following topics:

- [Configuring AAA Accounting for Roles](#)
- [Applying an Accounting Method to a Role](#)
- [Displaying Active Accounting Sessions for Roles](#)

Configuring AAA Accounting for Roles

To configure AAA accounting for roles, use the **aaa accounting** command in CONFIGURATION mode.

```
aaa accounting {system | exec | commands {level | role role-name}} {name |
default} {start-stop | wait-start | stop-only} {tacacs+}
```

Example of Configuring AAA Accounting for Roles

The following example shows you how to configure AAA accounting to monitor commands executed by the users who have a `secadmin` user role.

```
Dell(conf)#aaa accounting command role secadmin default start-stop tacacs+
```

Applying an Accounting Method to a Role

To apply an accounting method list to a role executed by a user with that user role, use the `accounting` command in LINE mode.

```
accounting {exec | commands {level | role role-name}} method-list
```

Example of Applying an Accounting Method to a Role

The following example applies the accounting default method to the user role `secadmin` (security administrator).

```
Dell(conf-vty-0)# accounting commands role secadmin default
```

Displaying Active Accounting Sessions for Roles

To display active accounting sessions for each user role, use the **show accounting** command in EXEC mode.

Example of Displaying Active Accounting Sessions for Roles

```
Dell#show accounting
```

```
Active accounted actions on tty2, User john Priv 1 Role netoperator
```

```
Task ID 1, EXEC Accounting record, 00:00:30 Elapsed,
```

```
service=shell
```

```
Active accounted actions on tty3, User admin Priv 15 Role sysadmin
```

```
Task ID 2, EXEC Accounting record, 00:00:26 Elapsed,
```

```
service=shell
```

Display Information About User Roles

This section describes how to display information about user roles.

This sections consists of the following topics:

- Displaying User Roles
- Displaying Information About Roles Logged into the Switch
- Displaying Active Accounting Sessions for Roles

Displaying User Roles

To display user roles using the `show userrole` command in EXEC Privilege mode, use the `show userroles` and `show users` commands in EXEC privilege mode.

Examples of Displaying User Roles

```
Dell#show userroles
Role          Inheritance      Modes
netoperator   Exec              Exec
netadmin      Exec Config Interface Line Router IP Routemap
Protocol MAC
secadmin      Exec Config
sysadmin      Exec Config Interface Line Router IP Routemap
Protocol MAC
testadmin     netadmin         Exec Config Interface Line Router IP Routemap
Protocol MAC
```

Displaying Role Permissions Assigned to a Command

To display permissions assigned to a command, use the `show role` command in EXEC Privilege mode. The output displays the user role and or permission level.

Examples of Role Permissions Assigned to a Command

```
Dell#show role mode ?
configure      Global configuration mode
exec           Exec Mode
interface      Interface configuration mode
line           Line Configuration mode
route-map      Route map configuration mode
router         Router configuration mode
```

```
Dell#show role mode configure username
Role access: sysadmin
```

```
Dell##show role mode configure password-attributes
Role access: secadmin,sysadmin
```

```
Dell#show role mode configure interface
Role access: netadmin, sysadmin
```

```
Dell#show role mode configure line
Role access: netadmin,sysadmin
```

Displaying Information About Users Logged into the Switch

To display information on all users logged into the switch, using the `show users` command in EXEC Privilege mode. The output displays privilege level and/or user role. The mode is displayed at the start of

the output and both the privilege and roles for all users is also displayed. If the role is not defined, the system displays "unassigned" .

Example of Displaying Information About Users Logged into the Switch

```
Dell#show users
```

```
Authorization Mode:  role or privilege
```

Line	User	Role	Privilege	Host(s)	Location
0 console 0	admin	sysadmin	15	idle	
*3 vty 1	secl	secadmin	14	idle	172.31.1.4
4 vty 2	ml1	netadmin	12	idle	172.31.1.5

Service Provider Bridging

Service provider bridging is supported on the MXL switch platform.

VLAN Stacking

VLAN stacking, also called Q-in-Q, is defined in IEEE 802.1ad — *Provider Bridges*, which is an amendment to IEEE 802.1Q — *Virtual Bridged Local Area Networks*.

VLAN stacking enables service providers to use 802.1Q architecture to offer separate VLANs to customers with no coordination between customers, and minimal coordination between customers and the provider.

Using only 802.1Q VLAN tagging all customers would have to use unique VLAN IDs to ensure that traffic is segregated, and customers and the service provider would have to coordinate to ensure that traffic mapped correctly across the provider network. Even under ideal conditions, customers and the provider would still share the 4094 available VLANs.

Instead, 802.1ad allows service providers to add their own VLAN tag to frames traversing the provider network. The provider can then differentiate customers even if they use the same VLAN ID, and providers can map multiple customers to a single VLAN to overcome the 4094 VLAN limitation. Forwarding decisions in the provider network are based on the provider VLAN tag only, so the provider can map traffic through the core independently; the customer and provider only coordinate at the provider edge.

At the access point of a VLAN-stacking network, service providers add a VLAN tag, the S-Tag, to each frame before the 802.1Q tag. From this point, the frame is double-tagged. The service provider uses the S-Tag, to forward the frame traffic across its network. At the egress edge, the provider removes the S-Tag, so that the customer receives the frame in its original condition, as shown in the following illustration.

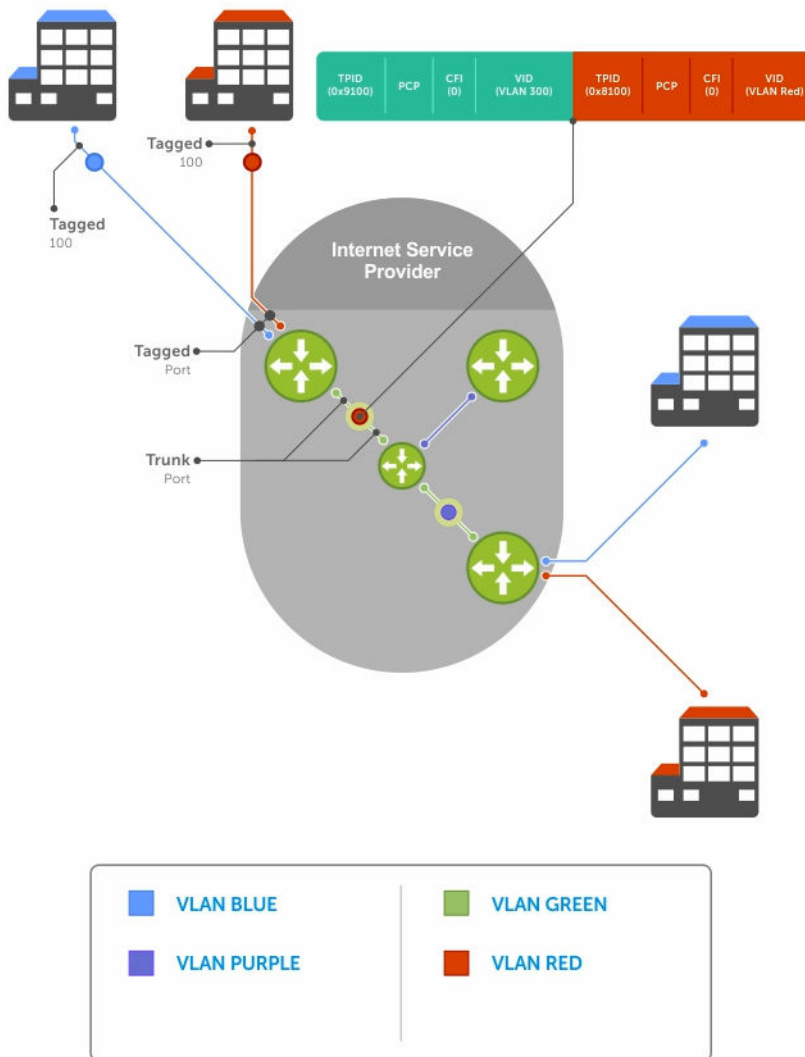


Figure 110. VLAN Stacking in a Service Provider Network

Important Points to Remember

- Interfaces that are members of the Default VLAN and are configured as VLAN-Stack access or trunk ports do not switch untagged traffic. To switch traffic, add these interfaces to a non-default VLAN-Stack-enabled VLAN.
- Dell Networking cautions against using the same MAC address on different customer VLANs, on the same VLAN-Stack VLAN.

Configure VLAN Stacking

Configuring VLAN-Stacking is a three-step process.

1. [Creating Access and Trunk Ports](#)
2. Assign access and trunk ports to a VLAN ([Creating Access and Trunk Ports](#)).
3. [Enable VLAN-Stacking for a VLAN](#).

Related Configuration Tasks

- [Configuring the Protocol Type Value for the Outer VLAN Tag](#)
- [Configuring Options for Trunk Ports](#)
- [Debugging VLAN Stacking](#)
- [VLAN Stacking in Multi-Vendor Networks](#)

Creating Access and Trunk Ports

To create access and trunk ports, use the following commands.

- **Access port** — a port on the service provider edge that directly connects to the customer. An access port may belong to only one service provider VLAN.
- **Trunk port** — a port on a service provider bridge that connects to another service provider bridge and is a member of multiple service provider VLANs.

Physical ports and port-channels can be access or trunk ports.

1. Assign the role of access port to a Layer 2 port on a provider bridge that is connected to a customer.
INTERFACE mode

```
vlan-stack access
```

2. Assign the role of trunk port to a Layer 2 port on a provider bridge that is connected to another provider bridge.
INTERFACE mode

```
vlan-stack trunk
```

3. Assign all access ports and trunk ports to service provider VLANs.
INTERFACE VLAN mode

```
member
```

Example of Displaying the VLAN-Stack Configuration for a Switchport

To display the VLAN-Stacking configuration for a switchport, use the `show config` command from INTERFACE mode.

```
Dell#show run interface gi 7/0
!
interface GigabitEthernet 7/0
  no ip address
  switchport
  vlan-stack access
  no shutdown
```

```
Dell#show run interface gi 7/12
!
interface GigabitEthernet 7/12
  no ip address
  switchport
  vlan-stack trunk
  no shutdown
```

Enable VLAN-Stacking for a VLAN

To enable VLAN-Stacking for a VLAN, use the following command.

- Enable VLAN-Stacking for the VLAN.

```
INTERFACE VLAN mode
```

```
  vlan-stack compatible
```

Example of Viewing VLAN Stack Member Status

To display the status and members of a VLAN, use the `show vlan` command from EXEC Privilege mode. Members of a VLAN-Stacking-enabled VLAN are marked with an M in column Q.

```
Dell#show vlan
```

```
Codes: * - Default VLAN, G - GVRP VLANs
```

```
NUM  Status      Q Ports
* 1  Active      U Gi 13/0-5,18
  2  Inactive
  3  Inactive
  4  Inactive
  5  Inactive
  6  Active      M Po1 (Gi 13/14-15)
                        M Gi 13/13
```

```
Dell#
```

Configuring the Protocol Type Value for the Outer VLAN Tag

The tag protocol identifier (TPID) field of the S-Tag is user-configurable.

To set the S-Tag TPID, use the following command.

- Select a value for the S-Tag TPID.

```
CONFIGURATION mode
```

```
  vlan-stack protocol-type
```

The default is **9100**.

To display the S-Tag TPID for a VLAN, use the `show running-config` command from EXEC privilege mode. The system displays the S-Tag TPID only if it is a non-default value.

Configuring Options for Trunk Ports

802.1ad trunk ports may also be tagged members of a VLAN so that it can carry single and double-tagged traffic.

You can enable trunk ports to carry untagged, single-tagged, and double-tagged VLAN traffic by making the trunk port a hybrid port.

To configure trunk ports, use the following commands.

1. Configure a trunk port to carry untagged, single-tagged, and double-tagged traffic by making it a hybrid port.

```
INTERFACE mode
```

```
portmode hybrid
```



NOTE: You can add a trunk port to an 802.1Q VLAN as well as a Stacking VLAN only when the TPID 0x8100.

2. Add the port to a 802.1Q VLAN as tagged or untagged.

```
INTERFACE VLAN mode
```

```
[tagged | untagged]
```

Example of Configuring a Trunk Port as a Hybrid Port and Adding it to Stacked VLANs

In the following example, GigabitEthernet 0/1 is a trunk port that is configured as a hybrid port and then added to VLAN 100 as untagged VLAN 101 as tagged, and VLAN 103, which is a stacking VLAN.

```
Dell(conf)#int gi 0/1
Dell(conf-if-gi-0/1)#portmode hybrid
Dell(conf-if-gi-0/1)#switchport
Dell(conf-if-gi-0/1)#vlan-stack trunk
Dell(conf-if-gi-0/1)#show config
!
interface GigabitEthernet 0/1
  no ip address
  portmode hybrid
  switchport
  vlan-stack trunk
  shutdown
Dell(conf-if-gi-0/1)#interface vlan 100
Dell(conf-if-vl-100)#untagged gigabitethernet 0/1
Dell(conf-if-vl-100)#interface vlan 101
Dell(conf-if-vl-101)#tagged gigabitethernet 0/1
Dell(conf-if-vl-101)#interface vlan 103
Dell(conf-if-vl-103)#vlan-stack compatible
Dell(conf-if-vl-103-stack)#member gigabitethernet 0/1
Dell(conf-if-vl-103-stack)#do show vlan
```

Codes: * - Default VLAN, G - GVRP VLANs

Q: U - Untagged, T - Tagged

x - Dot1x untagged, X - Dot1x tagged

G - GVRP tagged, M - Vlan-stack

NUM	Status	Description	Q	Ports
* 1	Inactive			
100	Inactive		U	Gi 0/1
101	Inactive		T	Gi 0/1
103	Inactive		M	Gi 0/1

Debugging VLAN Stacking

To debug VLAN stacking, use the following command.

- Debug the internal state and membership of a VLAN and its ports.

```
debug member
```

Example of Debugging a VLAN and its Ports

The port notations are as follows:

- **MT** – stacked trunk
- **MU** – stacked access port
- **T** – 802.1Q trunk port
- **U** – 802.1Q access port
- **NU** – Native VLAN (untagged)

```
Dell# debug member vlan 603
vlan id   : 603
ports    : Gi 2/47 (MT), Gi 3/1 (MU), Gi 3/25 (MT), Gi 3/26 (MT), Gi 3/27 (MU)
```

```
Dell# debug member port gigabitethernet 2/47
vlan id   : 603 (MT), 100 (T), 101 (NU)
Dell#
```

VLAN Stacking in Multi-Vendor Networks

The first field in the VLAN tag is the tag protocol identifier (TPID), which is 2 bytes. In a VLAN-stacking network, after the frame is double tagged, the outer tag TPID must match the TPID of the next-hop system.

While 802.1Q requires that the inner tag TPID is 0x8100, it does not require a specific value for the outer tag TPID. Systems may use any 2-byte value; the Dell Networking OS uses 0x9100 (shown in the following) while non-Dell Networking systems might use a different value.

If the next-hop system's TPID does not match the outer-tag TPID of the incoming frame, the system drops the frame. For example, as shown in the following, the frame originating from Building A is tagged VLAN RED, and then double-tagged VLAN PURPLE on egress at R4. The TPID on the outer tag is 0x9100. R2's TPID must also be 0x9100, and it is, so R2 forwards the frame.

Given the matching-TPID requirement, there are limitations when you employ Dell Networking systems at network edges, at which, frames are either double tagged on ingress (R4) or the outer tag is removed on egress (R3).

VLAN Stacking

The default TPID for the outer VLAN tag is 0x9100. Beginning with the Dell Networking OS version 8.2.1.0, the system allows you to configure both bytes of the 2 byte TPID.

Previous versions allowed you to configure the first byte only, and thus, the systems did not differentiate between TPIDs with a common first byte. For example, 0x8100 and any other TPID beginning with 0x81 were treated as the same TPID, as shown in the following illustration. The Dell Networking OS Versions 8.2.1.0 and later differentiate between 0x9100 and 0x91XY, also shown in the following illustration.

You can configure the first 8 bits of the TPID using the `vlan-stack protocol-type` command.

The TPID is global. Ingress frames that do not match the system TPID are treated as untagged. This rule applies for both the outer tag TPID of a double-tagged frame and the TPID of a single-tagged frame.

For example, if you configure TPID 0x9100, the system treats 0x8100 and untagged traffic the same and maps both types to the default VLAN, as shown by the frame originating from Building C. For the same traffic types, if you configure TPID 0x8100, the system is able to differentiate between 0x8100 and untagged traffic and maps each to the appropriate VLAN, as shown by the packet originating from Building A.

Therefore, a mismatched TPID results in the port not differentiating between tagged and untagged traffic.

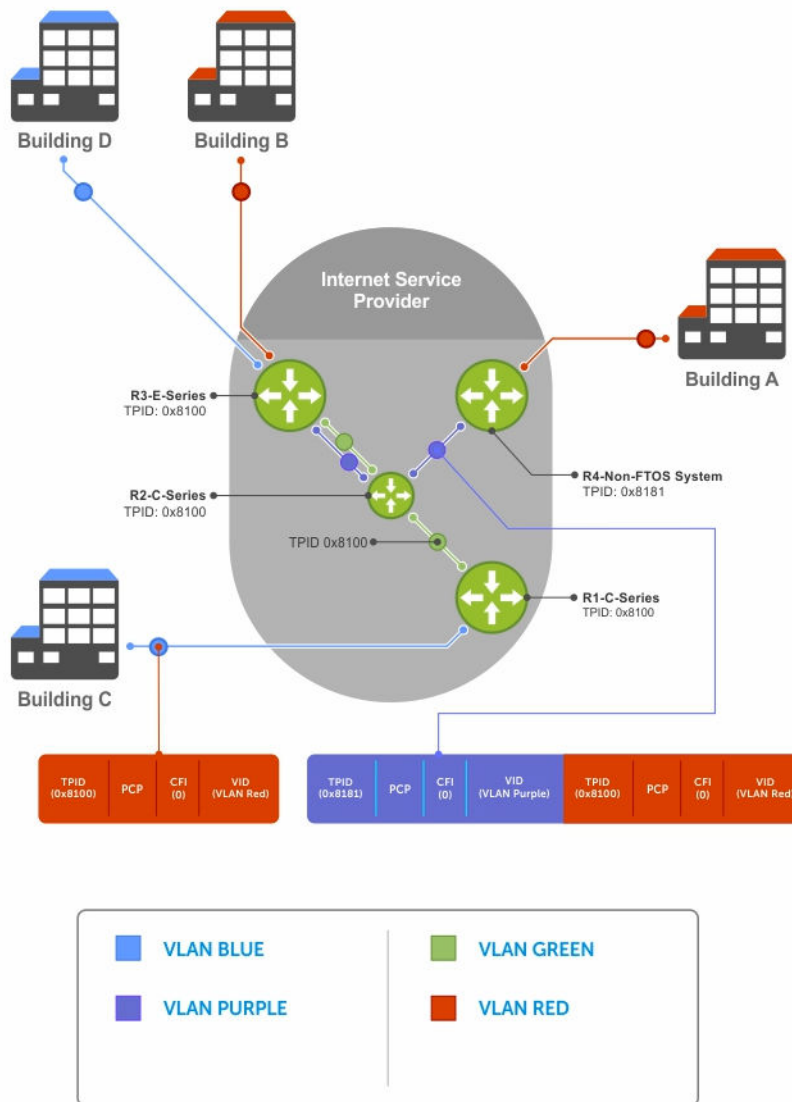


Figure 111. Single and Double-Tag TPID Match




Figure 112. Single and Double-Tag First-byte TPID Match

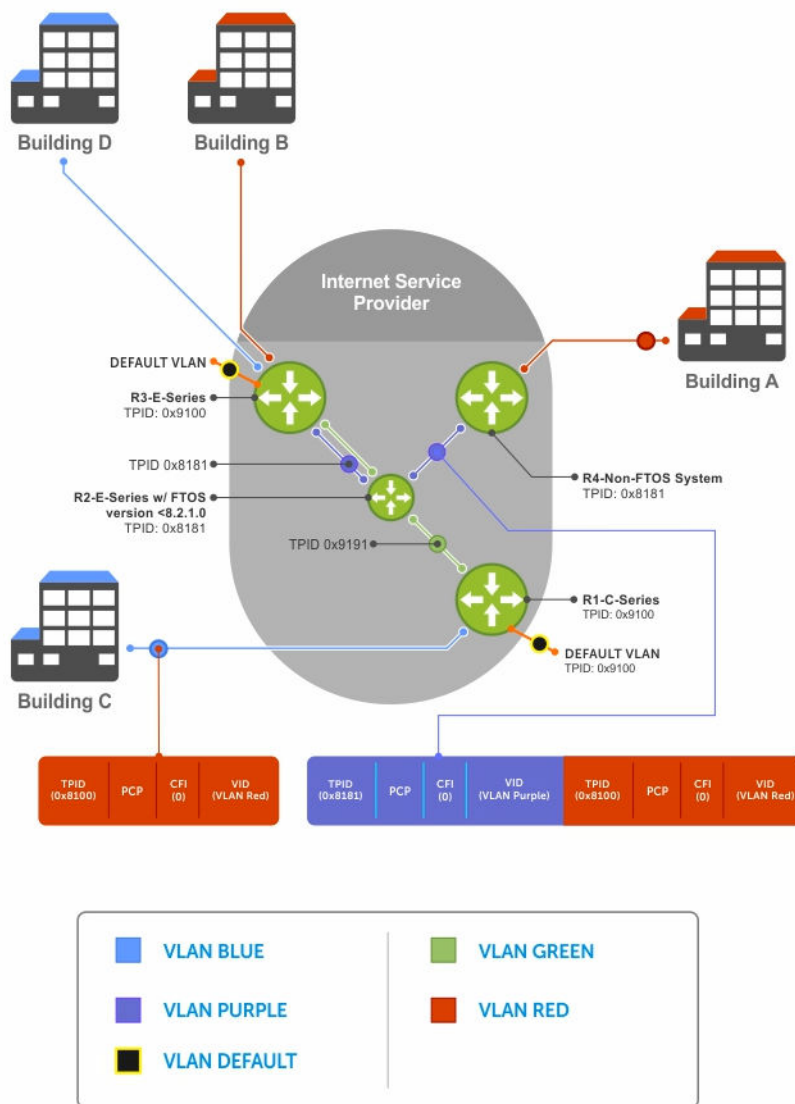


Figure 113. Single and Double-Tag TPID Mismatch

Table 57. Behaviors for Mismatched TPID

Network Position	Incoming Packet TPID	System TPID	Match Type	Pre-Version 8.2.1.0	Version 8.2.1.0+
Ingress Access Point	untagged	0xUVWX	—	switch to default VLAN	switch to default VLAN
	single-tag (0x8100)	0xUVWX	single-tag mismatch	switch to default VLAN	switch to default VLAN
		0x8100	single-tag match	switch to VLAN	switch to VLAN

Network Position	Incoming Packet TPID	System TPID	Match Type	Pre-Version 8.2.1.0	Version 8.2.1.0+
Core	untagged	0x81XY	single-tag first-byte match	switch to VLAN	switch to default VLAN
		0xUVWX	—	switch to default VLAN	switch to default VLAN
		0xUVWX	double-tag match	switch to VLAN	switch to VLAN
		0xUVYZ	double-tag first-byte match	switch to VLAN	switch to default VLAN
Egress Access Point	untagged	0xQRST	double-tag mismatch	switch to default VLAN	switch to default VLAN
		0xUVWX	—	switch to default VLAN	switch to default VLAN
		0xUVWX	double-tag match	switch to VLAN	switch to VLAN
		0xUVYZ	double-tag first-byte match	switch to VLAN	switch to default VLAN
		0xQRST	double-tag mismatch	switch to default VLAN	switch to default VLAN

VLAN Stacking Packet Drop Precedence

The drop eligible indicator (DEI) bit in the S-Tag indicates to a service provider bridge which packets it should prefer to drop when congested.

Enabling Drop Eligibility

Enable drop eligibility globally before you can honor or mark the DEI value. When you enable drop eligibility, DEI mapping or marking takes place according to the defaults. In this case, the CFI is affected according to the following table.

Table 58. Drop Eligibility Behavior

Ingress	Egress	DEI Disabled	DEI Enabled
Normal Port	Normal Port	Retain CFI	Set CFI to 0.
Trunk Port	Trunk Port	Retain inner tag CFI	Retain inner tag CFI.
		Retain outer tag CFI	Set outer tag CFI to 0.
Access Port	Trunk Port	Retain inner tag CFI	Retain inner tag CFI
		Set outer tag CFI to 0	Set outer tag CFI to 0

To enable drop eligibility globally, use the following command.

- Make packets eligible for dropping based on their DEI value.

CONFIGURATION mode

```
dei enable
```

By default, packets are colored green, and DEI is marked 0 on egress.

Honoring the Incoming DEI Value

To honor the incoming DEI value, you must explicitly map the DEI bit to a Dell Networking OS drop precedence.

Precedence can have one of three colors.

Precedence	Description
Green	High-priority packets that are the least preferred to be dropped.
Yellow	Lower-priority packets that are treated as best-effort.
Red	Lowest-priority packets that are always dropped (regardless of congestion status).

- Honor the incoming DEI value by mapping it to the Dell Networking OS drop precedence.

INTERFACE mode

```
dei honor {0 | 1} {green | red | yellow}
```

You may enter the command once for 0 and once for 1.

Packets with an unmapped DEI value are colored green.

Example of Viewing DEI-Honoring Configuration

To display the DEI-honoring configuration, use the `show interface dei-honor [interface slot/port | linecard number port-set number]` in EXEC Privilege mode.

```
Dell#show interface dei-honor
```

```
Default Drop precedence: Green
```

```
Interface CFI/DEI Drop precedence
```

```
-----
Gi 0/1      0      Green
Gi 0/1      1      Yellow
Gi 8/9      1      Red
Gi 8/40     0      Yellow
```

Marking Egress Packets with a DEI Value

On egress, you can set the DEI value according to a different mapping than ingress.

For ingress information, refer to [Honoring the Incoming DEI Value](#).

To mark egress packets, use the following command.

- Set the DEI value on egress according to the color currently assigned to the packet.

INTERFACE mode

```
dei mark {green | yellow} {0 | 1}
```

Example of Viewing DEI-Marking Configuration

To display the DEI-marking configuration, use the `show interface dei-mark [interface slot/port | linecard number port-set number]` in EXEC Privilege mode.

```
Dell#show interface dei-mark

Default CFI/DEI Marking: 0
Interface Drop precedence CFI/DEI
-----
Gi 0/1    Green    0
Gi 0/1    Yellow   1
Gi 8/9    Yellow   0
Gi 8/40   Yellow   0
```

Dynamic Mode CoS for VLAN Stacking

One of the ways to ensure quality of service for customer VLAN-tagged frames is to use the 802.1p priority bits in the tag to indicate the level of QoS desired.

When an S-Tag is added to incoming customer frames, the 802.1p bits on the S-Tag may be configured statically for each customer or derived from the C-Tag using Dynamic Mode CoS. Dynamic Mode CoS maps the C-Tag 802.1p value to a S-Tag 802.1p value.

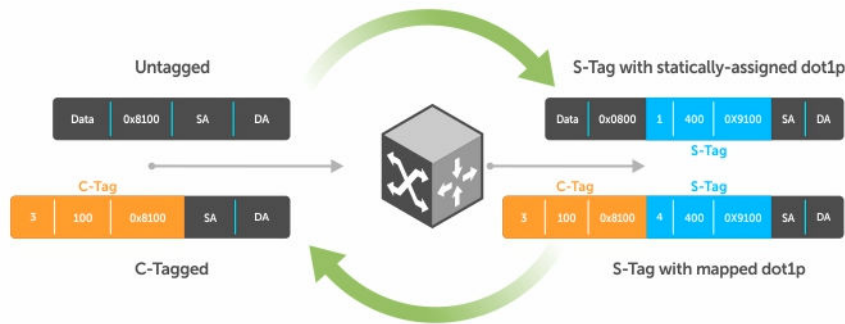


Figure 114. Statically and Dynamically Assigned dot1p for VLAN Stacking

When configuring Dynamic Mode CoS, you have two options:

- Mark the S-Tag dot1p and queue the frame according to the original C-Tag dot1p. In this case, you must have other dot1p QoS configurations; this option is classic dot1p marking.
- Mark the S-Tag dot1p and queue the frame according to the S-Tag dot1p. For example, if frames with C-Tag dot1p values 0, 6, and 7 are mapped to an S-Tag dot1p value 0, all such frames are sent to the queue associated with the S-Tag 802.1p value 0. This option requires two different CAM entries, each in a different Layer 2 ACL FP block.

NOTE: The ability to map incoming C-Tag dot1p to any S-Tag dot1p requires installing up to eight entries in the Layer 2 QoS and Layer 2 ACL table for each configured customer VLAN. The scalability of this feature is limited by the impact of the 1:8 expansion in these content addressable memory (CAM) tables.

Dell Networking OS Behavior: For Option A shown in the previous illustration, when there is a conflict between the queue selected by Dynamic Mode CoS (vlan-stack dot1p-mapping) and a QoS configuration, the queue selected by Dynamic Mode CoS takes precedence. However, rate policing for the queue is determined by QoS configuration. For example, the following access-port configuration maps all traffic to Queue 0:

```
vlan-stack dot1p-mapping c-tag-dot1p 0-7 sp-tag-dot1p 1
```

However, if the following QoS configuration also exists on the interface, traffic is queued to Queue 0 but is policed at 40Mbps (qos-policy-input for queue 3) because class-map "a" of Queue 3 also matches the traffic. This is an expected behavior.

Examples of QoS Interface Configuration and Rate Policing

```
policy-map-input in layer2
service-queue 3 class-map a qos-policy 3
!
class-map match-any a layer2
match mac access-group a
!
mac access-list standard a
seq 5 permit any
!
qos-policy-input 3 layer2
rate-police 40
```

Likewise, in the following configuration, packets with dot1p priority 0–3 are marked as dot1p 7 in the outer tag and queued to Queue 3. Rate policing is according to qos-policy-input 3. All other packets will have outer dot1p 0 and hence are queued to Queue 1. They are therefore policed according to qos-policy-input 1.

```
policy-map-input in layer2
  service-queue 1 qos-policy 1
  service-queue 3 qos-policy 3
!
qos-policy-input 1 layer2
  rate-police 10
!
qos-policy-input 3 layer2
  rate-police 30
!
interface GigabitEthernet 0/21
  no ip address
  switchport
  vlan-stack access
  vlan-stack dot1p-mapping c-tag-dot1p 0-3 sp-tag-dot1p 7
  service-policy input in layer2
  no shutdown
```

Mapping C-Tag to S-Tag dot1p Values

To map C-Tag dot1p values to S-Tag dot1p values and mark the frames accordingly, use the following commands.

1. Allocate CAM space to enable queuing frames according to the C-Tag or the S-Tag.

CONFIGURATION mode

```
cam-acl l2acl number ipv4acl number ipv6acl number ipv4qos number l2qos
number l2pt number ipmacacl number ecfmacacl number {vman-qos | vman-qos-dual-
fp} number
```

- `vman-qos`: mark the S-Tag dot1p and queue the frame according to the original C-Tag dot1p. This method requires half as many CAM entries as `vman-qos-dual-fp`.
- `vman-qos-dual-fp`: mark the S-Tag dot1p and queue the frame according to the S-Tag dot1p. This method requires twice as many CAM entries as `vman-qos` and FP blocks in multiples of 2.

The default is: 0 FP blocks for `vman-qos` and `vman-qos-dual-fp`.

2. The new CAM configuration is stored in NVRAM and takes effect only after a save and reload.
EXEC Privilege mode

```
copy running-config startup-config reload
```


3. Map C-Tag dot1p values to a S-Tag dot1p value.

INTERFACE mode

```
vlan-stack dot1p-mapping c-tag-dot1p values sp-tag-dot1p value
```

Separate C-Tag values by commas. Dashed ranges are permitted.

Dynamic Mode CoS overrides any Layer 2 QoS configuration in case of conflicts.

 **NOTE:** Because `dot1p-mapping` marks *and* queues packets, the only remaining applicable QoS configuration is rate metering. You may use Rate Shaping or Rate Policing.

Layer 2 Protocol Tunneling

Spanning tree bridge protocol data units (BPDUs) use a reserved destination MAC address called the bridge group address, which is 01-80-C2-00-00-00.

Only spanning-tree bridges on the local area network (LAN) recognize this address and process the BPDU. When you use VLAN stacking to connect physically separate regions of a network, BPDUs attempting to traverse the intermediate network might be consumed and later dropped because the intermediate network itself might be using spanning tree (shown in the following illustration).




Figure 115. VLAN Stacking without L2PT

You might need to transport control traffic transparently through the intermediate network to the other region. Layer 2 protocol tunneling enables BPDUs to traverse the intermediate network by identifying frames with the Bridge Group Address, rewriting the destination MAC to a user-configured non-reserved address, and forwarding the frames. Because the frames now use a unique MAC address, BPDUs are treated as normal data frames by the switches in the intermediate network core. On egress edge of the intermediate network, the MAC address rewritten to the original MAC address and forwarded to the opposing network region (shown in the following illustration).

Dell Networking OS Behavior: In the Dell Networking OS versions prior to 8.2.1.0, the MAC address that Dell Networking systems use to overwrite the Bridge Group Address on ingress was non-configurable. The value of the L2PT MAC address was the Dell Networking-unique MAC address, 01-01-e8-00-00-00. As such, with these versions, Dell Networking systems are required at the egress edge of the intermediate

network because only the Dell Networking OS could recognize the significance of the destination MAC address and rewrite it to the original Bridge Group Address. In the Dell Networking OS version 8.2.1.0 and later, the L2PT MAC address is user-configurable, so you can specify an address that non-Dell Networking systems can recognize and rewrite the address at egress edge.




Figure 116. VLAN Stacking with L2PT

Implementation Information

- L2PT is available for STP, RSTP, MSTP, and PVST+ BPDUs.
- No protocol packets are tunneled when you enable VLAN stacking.
- L2PT requires the default CAM profile.

Enabling Layer 2 Protocol Tunneling

To enable Layer 2 protocol tunneling, use the following command.

1. Verify that the system is running the default CAM profile. Use this CAM profile for L2PT.

EXEC Privilege mode

```
show cam-profile
```

2. Enable protocol tunneling globally on the system.

CONFIGURATION mode

```
protocol-tunnel enable
```

3. Tunnel BPDUs the VLAN.

INTERFACE VLAN mode

```
protocol-tunnel stp
```

Specifying a Destination MAC Address for BPDUs

By default, the Dell Networking OS uses a Dell Networking-unique MAC address for tunneling BPDUs. You can configure another value.

To specify a destination MAC address for BPDUs, use the following command.

- Overwrite the BPDU with a user-specified destination MAC address when BPDUs are tunneled across the provider network.

CONFIGURATION mode

```
protocol-tunnel destination-mac
```

The default is 01:01:e8:00:00:00

Setting Rate-Limit BPDUs

CAM space is allocated in sections called field processor (FP) blocks.

There are a total of 13 user-configurable FP blocks. The default number of blocks for L2PT is **0**; you must allocate at least one to enable BPDU rate-limiting.

To set the rate-lime BPDUs, use the following commands.

1. Create at least one FP group for L2PT.

CONFIGURATION mode

```
cam-acl l2acl
```

For details about this command, refer to [CAM Allocation](#).

2. Save the running-config to the startup-config.

EXEC Privilege mode

```
copy running-config startup-config
```

3. Reload the system.
EXEC Privilege mode

`reload`
4. Set a maximum rate at which the RPM processes BPDUs for L2PT.
VLAN STACKING mode

`protocol-tunnel rate-limit`

The default is: no rate limiting.

The range is from 64 to 320 kbps.

Debugging Layer 2 Protocol Tunneling

To debug Layer 2 protocol tunneling, use the following command.

- Display debugging information for L2PT.
EXEC Privilege mode

`debug protocol-tunnel`

Provider Backbone Bridging

IEEE 802.1ad—Provider Bridges amends 802.1Q—Virtual Bridged Local Area Networks so that service providers can use 802.1Q architecture to offer separate VLANs to customers with no coordination between customers, and minimal coordination between customers and the provider.

802.1ad specifies that provider bridges operating spanning tree use a reserved destination MAC address called the Provider Bridge Group Address, 01-80-C2-00-00-08, to exchange BPDUs instead of the Bridge Group Address, 01-80-C2-00-00-00, originally specified in 802.1Q. Only bridges in the service provider network use this destination MAC address so these bridges treat BPDUs originating from the customer network as normal data frames, rather than consuming them.

The same is true for GARP VLAN registration protocol (GVRP). 802.1ad specifies that provider bridges participating in GVRP use a reserved destination MAC address called the Provider Bridge GVRP Address, 01-80-C2-00-00-0D, to exchange GARP PDUs instead of the GVRP Address, 01-80-C2-00-00-21, specified in 802.1Q. Only bridges in the service provider network use this destination MAC address so these bridges treat GARP PDUs originating from the customer network as normal data frames, rather than consuming them.

Provider backbone bridging through IEEE 802.1ad eliminates the need for tunneling BPDUs with L2PT and increases the reliability of provider bridge networks as the network core need only learn the MAC addresses of core switches, as opposed to all MAC addresses received from attached customer devices.

sFlow

Configuring sFlow is supported on the MXL switch platform.

Overview

The Dell Networking operating system (OS) supports sFlow version 5.

sFlow is a standard-based sampling technology embedded within switches and routers which is used to monitor network traffic. It is designed to provide traffic monitoring for high-speed networks with many switches and routers. sFlow uses two types of sampling:

- Statistical packet-based sampling of switched or routed packet flows.
- Time-based sampling of interface counters.

The sFlow monitoring system consists of an sFlow agent (embedded in the switch/router) and an sFlow collector. The sFlow agent resides anywhere within the path of the packet and combines the flow samples and interface counters into sFlow datagrams and forwards them to the sFlow collector at regular intervals. The datagrams consist of information on, but not limited to, packet header, ingress and egress interfaces, sampling parameters, and interface counters.

Application-specific integrated circuits (ASICs) typically complete packet sampling. sFlow collector analyses the sFlow datagrams received from different devices and produces a network-wide view of traffic flows.

Implementation Information

Dell Networking sFlow is designed so that the hardware sampling rate is per line card port-pipe and is decided based on all the ports in that port-pipe.

If you do not enable sFlow on any port specifically, the global sampling rate is downloaded to that port and is to calculate the port-pipe's lowest sampling rate. This design supports the possibility that sFlow might be configured on that port in the future. Back-off is triggered based on the port-pipe's hardware sampling rate.

The default global sampling rate is 32768. The sampling rates are determined as follows:

- If the interface stats are up and the sampling rate is not configured on the port, the default sampling rate is calculated based on the line speed.
- If the interface states are shutdown, the sampling rate is set using the global sampling rate.
- If the global sample rate is non-default, for example 256 bytes, and if the sampling rate is not configured on an interface, the sampling rate of the interface is the global non-default sampling rate, that is 256 bytes.

To avoid the back-off, either increase the global sampling rate or configure all the line card ports with the desired sampling rate even if some ports have no sFlow configured.

Important Points to Remember

- The Dell Networking OS implementation of the sFlow MIB supports sFlow configuration using the `snmpset` command.
- The Dell Networking OS exports all sFlow packets to the collector. A small sampling rate can equate to many exported packets. A backoff mechanism is automatically applied to reduce this amount. Some sampled packets may be dropped when the exported packet rate is high and the backoff mechanism is about to or is starting to take effect. The `dropEvent` counter, in the sFlow packet, is always zero.
- Community list and local preference fields are not filled in extended gateway element in the sFlow datagram.
- 802.1P source priority field is not filled in extended switch element in sFlow datagram.
- Only Destination and Destination Peer AS number are packed in the `dst-as-path` field in extended gateway element.
- If the packet being sampled is redirected using policy-based routing (PBR), the sFlow datagram may contain incorrect extended gateway/router information.
- The source virtual local area network (VLAN) field in the extended switch element is not packed in case of routed packet.
- The destination VLAN field in the extended switch element is not packed in a Multicast packet.
- On the MXL switch, up to 700 packets can be sampled and processed per second.

Enabling and Disabling sFlow

By default, sFlow is disabled globally on the system. Use the following command to enable sFlow globally.

- Enable sFlow globally.
CONFIGURATION mode

```
[no] sflow enable
```

Enabling and Disabling sFlow on an Interface

By default, sFlow is disabled on all interfaces. This CLI is supported on physical ports and link aggregation group (LAG) ports. To enable sFlow on a specific interface, use the following command.

- Enable sFlow on an interface.
INTERFACE mode

```
[no] sflow enable
```

To disable sFlow on an interface, use the `no` version of this command.

Enabling sFlow Max-Header Size Extended

To configure the maximum header size of a packet to 256 bytes, use the following commands:

- Set the maximum header size of a packet.

```
CONFIGURATION mode
```

```
INTERFACE mode
```

```
sflow max-header-size extended
```

By default, the maximum header size of a packet is 128 bytes. If the traffic ingresses on an sFlow enabled interface, 256 bytes are copied.

- To reset the maximum header size of a packet, use the following command

```
[no] sflow max-header-size extended
```

- View the maximum header size of a packet.

```
show running-config sflow
```

Example of Viewing the Extended Max Header Size

```
Dell#show sflow interface tengigabitethernet 1/1
Te 1/1
sFlow type                :Ingress   <<<<
Configured sampling rate  :16384
Actual sampling rate      :16384
Counter polling interval  :20
Extended max header size :256
Samples rcvd from h/w     :0
```

Example of the show sflow command

The bold line shows the sFlow default maximum header size:

```
Dell#show sflow
sFlow services are enabled
Egress Management Interface sFlow services are disabled
Global default sampling rate: 32768
Global default counter polling interval: 20
Global default extended maximum header size: 128 bytes
Global extended information enabled: none
1 collectors configured
Collector IP addr: 100.1.1.1, Agent IP addr: 1.1.1.2, UDP port: 6343 VRF:
Default
0 UDP packets exported
0 UDP packets dropped
0 sFlow samples collected

stack-unit 0 Port set 0
  Te 1/1: configured rate 16384, actual rate 16384 <<< sampling rate based on
line speed if global sampling rate is default
Dell#
```

Example of Viewing the sFlow max-header-size extended on an Interface Mode

If you enable sFlow on an interface, the show output displays the following (shown in bold).

```
Dell(conf-if-te-1/10)#show sflow
sFlow services are enabled
Egress Management Interface sFlow services are disabled
Global default sampling rate: 32768
Global default counter polling interval: 86400
Global default extended maximum header size: 256 bytes
Global extended information enabled: none
1 collectors configured
Collector IP addr: 100.1.1.12, Agent IP addr: 100.1.1.1, UDP port: 6343 VRF:
Default
0 UDP packets exported
0 UDP packets dropped
0 sFlow samples collected
```

Example of the show running-config sflow Command

```
Dell#show running-config sflow
!
sflow collector 100.1.1.12 agent-addr 100.1.1.1
sflow enable
sflow max-header-size extended
```

```
Dell#show run int tengigabitEthernet 1/10
!
interface TenGigabitEthernet 1/10
no ip address
switchport
sflow ingress-enable
sflow max-header-size extended
no shutdown
```

sFlow Show Commands

The Dell Networking OS includes the following sFlow display commands.

- [Displaying Show sFlow Global](#)
- [Displaying Show sFlow on an Interface](#)
- [Displaying Show sFlow on a Stack Unit](#)

Displaying Show sFlow Global

To view sFlow statistics, use the following command.

- Display sFlow configuration information and statistics.
EXEC mode

```
show sflow
```

Example of Viewing sFlow Configuration (Global)

The first bold line indicates sFlow is globally enabled.

```
Dell#show sflow
sFlow services are enabled
```

```
Global default sampling rate: 32768
Global default counter polling interval: 20
1 collectors configured
Collector IP addr: 133.33.33.53, Agent IP addr: 133.33.33.116, UDP port: 6343
77 UDP packets exported
0 UDP packets dropped
165 sFlow samples collected
69 sFlow samples dropped due to sub-sampling
Dell#
```

Displaying Show sFlow on an Interface

To view sFlow information on a specific interface, use the following command.

- Display sFlow configuration information and statistics on a specific interface.

EXEC mode

```
show sflow interface interface-name
```

Example of Viewing sFlow Configuration (Interface)

```
Dell#show sflow interface gigabitethernet 1/16
Tengig 1/16
Configured sampling rate      :8192
Actual sampling rate         :8192
Sub-sampling rate            :2
Counter polling interval     :15
Samples rcvd from h/w        :33
Samples dropped for sub-sampling :6
```

Displaying Show sFlow on a Stack Unit

To view sFlow statistics on a specified stack unit, use the following command.

- Display sFlow configuration information and statistics on the specified interface.

EXEC mode

```
show sflow stack-unit unit-number
```

Example of Viewing sFlow Configuration (Stack Unit)

```
Dell#show sflow stack-unit 1
Stack-Unit 1
  Samples rcvd from h/w      :0
  Total UDP packets exported :0
  UDP packets exported via RPM :0
  UDP packets dropped        :0
Dell#
```

Configuring Specify Collectors

The `sflow collector` command allows identification of sFlow collectors to which sFlow datagrams are forwarded.

You can specify up to two sFlow collectors. If you specify two collectors, the samples are sent to both.

- Identify sFlow collectors to which sFlow datagrams are forwarded.

CONFIGURATION mode

```
sflow collector ip-address agent-addr ip-address [number [max-datagram-size  
number] ] | [max-datagram-size number ]
```

The default UDP port is **6343**.

The default max-datagram-size is **1400**.

Changing the Polling Intervals

The `sflow polling-interval` command configures the polling interval for an interface in the maximum number of seconds between successive samples of counters sent to the collector. This command changes the global default counter polling (20 seconds) interval. You can configure an interface to use a different polling interval.

To configure the polling intervals globally (in CONFIGURATION mode) or by interface (in INTERFACE mode), use the following command.

- Change the global default counter polling interval.

CONFIGURATION mode or INTERFACE mode

```
sflow polling-interval interval value
```

- *interval value*: in seconds.

The range is from 15 to 86400 seconds.

The default is **20 seconds**.

Changing the Sampling Rate

The `sflow sample-rate` command, when issued in CONFIGURATION mode, changes the default sampling rate.

By default, the sampling rate of an interface is set to the same value as the current global default sampling rate. If the value entered is not a correct power of 2, the command generates an error message with the previous and next power-of-2 value. Select one of these two numbers and re-enter the command. (For more information on values in power-of-2, refer to [Sub-Sampling](#).)

To change the sampling rate either globally or on an interface, use the following command.

- Change the global or interface sampling rate.

CONFIGURATION mode or INTERFACE mode

```
[no] sflow sample-rate sample-rate
```

sample-rate: The range is from 256 to 8388608 for the C-Series and S-Series. The range is from 2 to 8388608 for the E-Series.

The rate must be entered in factors of 2 (for example, 4096 or 8192).

Sub-Sampling


The sFlow sample rate is not the frequency of sampling, but the number of packets that are skipped before the next sample is taken.

Therefore, the sFlow agent uses sub-sampling to create multiple sampling rates per port-pipe. To achieve different sampling rates for different ports in a port-pipe, the sFlow agent takes the lowest numerical value of the sampling rate of all the ports within the port-pipe and configures all the ports to this value. The sFlow agent is then able to skip samples on ports where you require a larger sampling rate value.

Sampling rates are configurable in powers of two. This configuration allows the smallest sampling rate possible on the hardware and also allows all other sampling rates to be available through sub-sampling.

For example, if Tengig 1/0 and 1/1 are in a port-pipe, and they are configured with a sampling rate of 4096 on interface Tengig 1/0, and 8192 on Tengig 1/1, the sFlow agent does the following:

1. Configures the hardware to a sampling rate of 4096 for all ports with sFlow enabled on that port-pipe.
2. Configures interface Tengig 1/0 to a sub-sampling rate of 1 to achieve an actual rate of 4096.
3. Configures interface Tengig 1/1 to a sub-sampling rate of 2 to achieve an actual rate of 8192.

 **NOTE:** Sampling rate backoff can change the sampling rate value that is set in the hardware. This equation shows the relationship between actual sampling rate, sub-sampling rate, and the hardware sampling rate for an interface:

$$\text{Actual sampling rate} = \text{sub-sampling rate} * \text{hardware sampling rate}$$

Note the absence of a configured rate in the equation. That absence is because when the hardware sampling rate value on the port-pipe exceeds the configured sampling rate value for an interface, the actual rate changes to the hardware rate. The sub-sampling rate never goes below a value of one.

Back-Off Mechanism

If the sampling rate for an interface is set to a very low value, the CPU can get overloaded with flow samples under high-traffic conditions.

In such a scenario, a binary back-off mechanism gets triggered, which doubles the sampling-rate (halves the number of samples per second) for all interfaces. The backoff mechanism continues to double the sampling-rate until the CPU condition is cleared. This is as per sFlow version 5 draft. After the back-off changes the sample-rate, you must manually change the sampling rate to the desired value.

As a result of back-off, the actual sampling-rate of an interface may differ from its configured sampling rate. You can view the actual sampling-rate of the interface and the configured sample-rate by using the `show sflow` command.

sFlow on LAG ports

When a physical port becomes a member of a LAG, it inherits the sFlow configuration from the LAG port.

Enabling Extended sFlow

The MXL switch support `extended-switch` information processing only. Extended sFlow packs additional information in the sFlow datagram depending on the type of sampled packet. You can enable the following options:

- `extended-switch` — 802.1Q VLAN ID and 802.1p priority information.
- `extended-router` — Next-hop and source and destination mask length.
- `extended-gateway` — Source and destination AS number and the BGP next-hop.

- Enable extended sFlow.

```
sflow [extended-switch] [extended-router] [extended-gateway] enable
```

By default packing of any of the extended information in the datagram is disabled.

- Confirm that extended information packing is enabled.

```
show sflow
```

Example of Verifying Extended sFlow is Enabled

Example of Verifying Extended sFlow Disabled

The bold line shows that extended sFlow settings are enabled on all three types.

```
Dell#show sflow
sFlow services are enabled
Global default sampling rate: 4096
Global default counter polling interval: 15
Global extended information enabled: switch
1 collectors configured
Collector IP addr: 10.10.10.3, Agent IP addr: 10.10.0.0, UDP port: 6343
77 UDP packets exported
0 UDP packets dropped
165 sFlow samples collected
69 sFlow samples dropped due to sub-sampling
Stackunit 1 Port set 0 H/W sampling rate 8192
Tengig 1/16: configured rate 8192, actual rate 8192, sub-sampling rate 1
Tengig 1/17: configured rate 16384, actual rate 16384, sub-sampling rate 2
Stackunit 3 Port set 1 H/W sampling rate 16384
Tengig 3/40: configured rate 16384, actual rate 16384, sub-sampling rate 1
```

If you did not enable any extended information, the show output displays the following (shown in bold).

```
Dell#show sflow
sFlow services are disabled
Global default sampling rate: 32768
Global default counter polling interval: 20
Global extended information enabled: none
0 collectors configured
```


```
0 UDP packets exported
0 UDP packets dropped
0 sFlow samples collected
0 sFlow samples dropped due to sub-sampling
```

Simple Network Management Protocol (SNMP)

Simple network management protocol (SNMP) is supported on the MXL switch platform.

Network management stations use SNMP to retrieve or alter management data from network elements. A datum of management information is called a managed object; the value of a managed object can be static or variable. Network elements store managed objects in a database called a management information base (MIB).

MIBs are hierarchically structured and use object identifiers to address managed objects, but managed objects also have a textual name called an *object descriptor*.

 **NOTE:** On Dell Networking routers, standard and private SNMP management information bases (MIBs) are supported, including all *Get* and a limited number of *Set* operations (such as `set vlan` and `copy cmd`).


Implementation Information

The following describes SNMP implementation information.

- The Dell Networking OS supports SNMP version 1 as defined by RFC 1155, 1157, and 1212, SNMP version 2c as defined by RFC 1901, and SNMP version 3 as defined by RFC 2571.
- The Dell Networking OS supports up to 16 trap receivers.
- The Dell Networking OS implementation of the sFlow MIB supports sFlow configuration via SNMP sets.
- SNMP traps for the spanning tree protocol (STP) and multiple spanning tree protocol (MSTP) state changes are based on BRIDGE MIB (RFC 1483) for STP and IEEE 802.1 *draft ruzin-mstp-mib-02* for MSTP.
- All objects in the LLDP-EXT-DOT1-DCBX-MIB and IEEE8021-PFC-MIB tables are read-only.
- In the LLDP-EXT-DOT1-DCBX-MIB, `lldpXdot1dcbxRemoteData` tables are not supported.

Configuration Task List for SNMP

Configuring SNMP version 1 or version 2 requires a single step.

 **NOTE:** The configurations in this chapter use a UNIX environment with `net-snmp` version 5.4. This environment is only one of many RFC-compliant SNMP utilities you can use to manage your Dell Networking system using SNMP. Also, these configurations use SNMP version 2c.

- [Creating a Community](#)

Configuring SNMP version 3 requires configuring SNMP users in one of three methods. Refer to [Setting Up User-Based Security \(SNMPv3\)](#).

Related Configuration Tasks

- [Set up SNMP](#)
- [Setting Up User-Based Security \(SNMPv3\)](#)
- [Reading Managed Object Values](#)
- [Writing Managed Object Values](#)
- [Configuring Contact and Location Information using SNMP](#)
- [Subscribing to Managed Object Value Updates using SNMP](#)
- [Copying Configuration Files via SNMP](#)
- [Manage VLANs using SNMP](#)
- [Enabling and Disabling a Port using SNMP](#)
- [Fetch Dynamic MAC Entries using SNMP](#)
- [Deriving Interface Indices](#)
- [Monitor Port-Channels](#)
- [Troubleshooting SNMP Operation](#)

Important Points to Remember

- Typically, 5-second timeout and 3-second retry values on an SNMP server are sufficient for both LAN and WAN applications. If you experience a timeout with these values, increase the timeout value to greater than 3 seconds, and increase the retry value to greater than 2 seconds on your SNMP server.
- User ACLs override group ACLs.

SNMPv3 Compliance With FIPS

SNMPv3 is compliant with the Federal information processing standard (FIPS) cryptography standard. The Advanced Encryption Standard (AES) Cipher Feedback (CFB) 128-bit encryption algorithm is in compliance with RFC 3826. SNMPv3 provides multiple authentication and privacy options for user configuration. A subset of these options are the FIPS-approved algorithms: HMAC-SHA1-96 for authentication and AES128-CFB for privacy. The other options are not FIPS-approved algorithms because of known security weaknesses. The AES128-CFB privacy option is supported and is compliant with RFC 3826.

The SNMPv3 feature also uses a FIPS-validated cryptographic module for all of its cryptographic operations when the system is configured with the `fips mode enable` command in Global Configuration mode. When the FIPS mode is enabled on the system, SNMPv3 operates in a FIPS-compliant manner, and only the FIPS-approved algorithm options are available for SNMPv3 user configuration. When the FIPS mode is disabled on the system, all options are available for SNMPv3 user configuration.

The following table describes the authentication and privacy options that can be configured when the FIPS mode is enabled or disabled:

FIPS Mode	Privacy Options	Authentication Options
Disabled	des56 (DES56-CBC)	md5 (HMAC-MD5-96)
	aes128 (AES128-CFB)	sha (HMAC-SHA1-96)
Enabled	aes128 (AES128-CFB)	sha (HMAC-SHA1-96)

To enable security for SNMP packets transferred between the server and the client, you can use the `snmp-server user username group groupname 3 auth authentication-type auth-password priv aes128 priv-password` command to specify that AES-CFB 128 encryption algorithm needs to be used.

```
Dell(conf)#snmp-server user snmpguy snmpmon 3 auth sha AArt61wq priv aes128
jntRR59a
```

In this example, for a specified user and a group, the AES128-CFB algorithm, the authentication password to enable the server to receive packets from the host, and the privacy password to encode the message contents are configured.

SHA authentication needs to be used with the AES-CFB128 privacy algorithm only when FIPS is enabled because SHA is then the only available authentication level. If FIPS is disabled, you can use MD5 authentication in addition to SHA authentication with the AES-CFB128 privacy algorithm

You cannot modify the FIPS mode if SNMPv3 users are already configured and present in the system. An error message is displayed if you attempt to change the FIPS mode by using the `fips mode enable` command in Global Configuration mode. You can enable or disable FIPS mode only if SNMPv3 users are not previously set up. If previously configured users exist on the system, you must delete the existing users before you change the FIPS mode.

Keep the following points in mind when you configure the AES128-CFB algorithm for SNMPv3:

1. SNMPv3 authentication provides only the `sha` option when the FIPS mode is enabled.
2. SNMPv3 privacy provides only the `aes128 privacy` option when the FIPS mode is enabled.
3. If you attempt to enable or disable FIPS mode and if any SNMPv3 users are previously configured, an error message is displayed stating you must delete all of the SNMP users before changing the FIPS mode.
4. A message is logged indicating whether FIPS mode is enabled for SNMPv3. This message is generated only when the first SNMPv3 user is configured because you can modify the FIPS mode only when users are not previously configured. This log message is provided to assist your system security auditing procedures.

Set up SNMP

As previously stated, the Dell Networking OS supports SNMP version 1 and version 2 that are community-based security models.

The primary difference between the two versions is that version 2 supports two additional protocol operations (*informs operation* and *snmpgetbulk query*) and one additional object (*counter64 object*).

SNMP version 3 (SNMPv3) is a user-based security model that provides password authentication for user security and encryption for data security and privacy. Three sets of configurations are available for SNMP read/write operations: no password or privacy, password privileges, password and privacy privileges.

You can configure a maximum of 16 users even if they are in different groups.

Creating a Community

For SNMPv1 and SNMPv2, create a community to enable the community-based security in the Dell Networking OS.

The management station generates requests to either retrieve or alter the value of a management object and is called the *SNMP manager*. A network element that processes SNMP requests is called an *SNMP agent*. An *SNMP community* is a group of SNMP agents and managers that are allowed to interact. Communities are necessary to secure communication between SNMP managers and agents; SNMP agents do not respond to requests from management stations that are not part of the community.

The Dell Networking OS enables SNMP automatically when you create an SNMP community and displays the following message. You must specify whether members of the community may only retrieve values (read), or retrieve and alter values (read-write).

```
22:31:23: %RPM1-P:CP %SNMP-6-SNMP_WARM_START: Agent Initialized - SNMP
WARM_START.
```

To choose a name for the community you create, use the following command.

- Choose a name for the community.

CONFIGURATION mode

```
snmp-server community name {ro | rw}
```

Example of Creating an SNMP Community

To view your SNMP configuration, use the `show running-config snmp` command from EXEC Privilege mode.

```
Dell(conf)#snmp-server community my-snmp-community ro
22:31:23: %RPM1-P:CP %SNMP-6-SNMP_WARM_START: Agent Initialized - SNMP
WARM_START.
Dell#show running-config snmp
!
snmp-server community mycommunity ro
Dell#
```

Setting Up User-Based Security (SNMPv3)

When setting up SNMPv3, you can set users up with one of the following three types of configuration for SNMP read/write operations.

Users are typically associated to an SNMP group with permissions provided, such as OID view.

- **noauth** — no password or privacy. Select this option to set up a user with no password or privacy privileges. This setting is the basic configuration. Users must have a group and profile that do not require password privileges.
- **auth** — password privileges. Select this option to set up a user with password authentication.
- **priv** — password and privacy privileges. Select this option to set up a user with password and privacy privileges.

To set up user-based security (SNMPv3), use the following commands.

- Configure the user with view privileges only (no password or privacy privileges).

CONFIGURATION mode

```
snmp-server user name group-name 3 noauth
```

- Configure an SNMP group with view privileges only (no password or privacy privileges).


CONFIGURATION mode

```
snmp-server group group-name 3 noauth auth read name write name
```

- Configure an SNMPv3 view.

CONFIGURATION mode

```
snmp-server view view-name oid-tree {included | excluded}
```

 **NOTE:** To give a user read and write view privileges, repeat this step for each privilege type.

- Configure the user with an authorization password (password privileges only).

CONFIGURATION mode

```
snmp-server user name group-name 3 noauth auth md5 auth-password
```

- Configure an SNMP group (password privileges only).

CONFIGURATION mode

```
snmp-server group groupname {oid-tree} auth read name write name
```

- Configure an SNMPv3 view.

CONFIGURATION mode

```
snmp-server view view-name 3 noauth {included | excluded}
```

 **NOTE:** To give a user read and write privileges, repeat this step for each privilege type.

- Configure an SNMP group (with password or privacy privileges).

CONFIGURATION mode

```
snmp-server group group-name {oid-tree} priv read name write name
```

- Configure the user with a secure authorization password and privacy password.

CONFIGURATION mode

```
snmp-server user name group-name {oid-tree} auth md5 auth-password priv des56  
priv password
```

- Configure an SNMPv3 view.

CONFIGURATION mode

```
snmp-server view view-name oid-tree {included | excluded}
```

Select a User-based Security Type

```
Dell(conf)#snmp-server host 1.1.1.1 traps {oid tree} version 3 ?  
auth          Use the SNMPv3 authNoPriv Security Level  
noauth       Use the SNMPv3 noAuthNoPriv Security Level  
priv         Use the SNMPv3 authPriv Security Level  
Dell(conf)#snmp-server host 1.1.1.1 traps {oid tree} version 3 noauth ?  
WORD          SNMPv3 user name
```

Reading Managed Object Values

You may only retrieve (read) managed object values if your management station is a member of the same community as the SNMP agent.

Dell Networking supports RFC 4001, *Textual Conventions for Internet Work Addresses* that defines values representing a type of internet address. These values display for `ipAddressTable` objects using the `snmpwalk` command.

There are several UNIX SNMP commands that read data.

- Read the value of a single managed object.

```
snmpget -v version -c community agent-ip {identifier.instance |
descriptor.instance}
```

- Read the value of the managed object directly below the specified object.

```
snmpgetnext -v version -c community agent-ip {identifier.instance |
descriptor.instance}
```

- Read the value of many objects at once.

```
snmpwalk -v version -c community agent-ip {identifier.instance |
descriptor.instance}
```

Example of Reading the Value of a Managed Object

Example of Reading the Value of the Next Managed Object

Example of Reading the Value of Many Managed Objects at Once

Example of the OID Display

In the following example, the value "4" displays in the OID before the IP address for IPv4. For an IPv6 IP address, a value of "16" displays.

```
> snmpget -v 2c -c mycommunity 10.11.131.161 sysUpTime.0
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (32852616) 3 days, 19:15:26.16
> snmpget -v 2c -c mycommunity 10.11.131.161 .1.3.6.1.2.1.1.3.0
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (32856932) 3 days, 19:16:09.32
```

```
> snmpgetnext -v 2c -c mycommunity 10.11.131.161 .1.3.6.1.2.1.1.3.0
SNMPv2-MIB::sysContact.0 = STRING:
> snmpgetnext -v 2c -c mycommunity 10.11.131.161 sysContact.0
SNMPv2-MIB::sysName.0 = STRING:
```

```
> snmpwalk -v 2c -c mycommunity 10.11.209.217 .1.3.6.1.2.1.1
SNMPv2-MIB::sysDescr.0 = STRING: Dell Networking OS
Operating System Version: 1.0
Application Software Version: E8-3-16-0
Series: MXL-10/40GbE
Copyright (c) 1999-2012 by Dell Inc. All Rights Reserved.
Build Time: Tue May 22 22:40:56 PDT 2012
SNMPv2-MIB::sysObjectID.0 = OID: SNMPv2-SMI::enterprises.6027.1.4.1
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (256676) 0:42:46.76
SNMPv2-MIB::sysContact.0 = STRING:
SNMPv2-MIB::sysName.0 = STRING: FTOS
SNMPv2-MIB::sysLocation.0 = STRING:
SNMPv2-MIB::sysServices.0 = INTEGER: 4
```

In the following example, the value **4** displays in the OID before the IP address for IPv4.

```
>snmpwalk -v 2c -c public 10.11.195.63 1.3.6.1.2.1.4.34
IP-MIB::ip.34.1.3.1.4.1.1.1.1 = INTEGER: 1107787778
IP-MIB::ip.34.1.3.1.4.2.1.1.1 = INTEGER: 1107787779
IP-MIB::ip.34.1.3.2.16.254.128.0.0.0.0.0.2.1.232.255.254.139.5.8 = INTEGER:
1107787778
IP-MIB::ip.34.1.4.1.4.1.1.1.1 = INTEGER: 1
IP-MIB::ip.34.1.4.1.4.2.1.1.1 = INTEGER: 1
IP-MIB::ip.34.1.4.2.16.0.1.0.0.0.0.0.0.0.0.0.0.0.0.1 = INTEGER: 1
```

Writing Managed Object Values

You may only alter (write) a managed object value if your management station is a member of the same community as the SNMP agent, and the object is writable.

Use the following command to write or write-over the value of a managed object.

- To write or write-over the value of a managed object.

```
snmpset -v version -c community agent-ip {identifier.instance |
descriptor.instance}syntax value
```

Example of Writing the Value of a Managed Object

```
> snmpset -v 2c -c mycommunity 10.11.131.161 sysName.0 s "R5"
SNMPv2-MIB::sysName.0 = STRING: R5
```

Configuring Contact and Location Information using SNMP

You may configure system contact and location information from the Dell Networking system or from the management station using SNMP.

To configure system contact and location information from the Dell Networking system and from the management station using SNMP, use the following commands.

- (From a Dell Networking system) Identify the system manager along with this person's contact information (for example, an email address or phone number).

CONFIGURATION mode

```
snmp-server contact text
```

You may use up to 55 characters.

The default is **None**.

- (From a Dell Networking system) Identify the physical location of the system (for example, San Jose, 350 Holger Way, 1st floor lab, rack A1-1).

CONFIGURATION mode

```
snmp-server location text
```

You may use up to 55 characters.

The default is **None**.

- (From a management station) Identify the system manager along with this person's contact information (for example, an email address or phone number).

CONFIGURATION mode

```
snmpset -v version -c community agent-ip sysContact.0 s "contact-info"
```

You may use up to 55 characters.

The default is **None**.

- (From a management station) Identify the physical location of the system (for example, San Jose, 350 Holger Way, 1st floor lab, rack A1-1).

CONFIGURATION mode

```
snmpset -v version -c community agent-ip sysLocation.0 s "location-info"
```

You may use up to 55 characters.

The default is **None**.

Subscribing to Managed Object Value Updates using SNMP

By default, the Dell Networking system displays some unsolicited SNMP messages (traps) upon certain events and conditions.

You can also configure the system to send the traps to a management station. Traps cannot be saved on the system.

The Dell Networking OS supports the following three sets of traps:

- **RFC 1157-defined traps** — coldStart, warmStart, linkDown, linkUp, authenticationFailure, and egpNeighborLoss.
- **Dell Networking enterpriseSpecific environment traps** — fan, supply, and temperature.
- **Dell Networking enterpriseSpecific protocol traps** — bgp, ecfm, stp, and xstp.

To configure the system to send SNMP notifications, use the following commands.

1. Configure the Dell Networking system to send notifications to an SNMP server.

CONFIGURATION mode

```
snmp-server host ip-address [traps | informs] [version 1 | 2c | 3]
[community-string]
```

To send trap messages, enter the keyword `traps`.

To send informational messages, enter the keyword `informs`.

To send the SNMP version to use for notification messages, enter the keyword `version`.

To identify the SNMPv1 community string, enter the name of the `community-string`.

2. Specify which traps the Dell Networking system sends to the trap receiver.

CONFIGURATION mode

```
snmp-server enable traps
```

Enable all Dell Networking enterprise-specific and RFC-defined traps using the `snmp-server enable traps` command from CONFIGURATION mode.

Enable all of the RFC-defined traps using the `snmp-server enable traps snmp` command from CONFIGURATION mode.

3. Specify the interfaces out of which the Dell Networking OS sends SNMP traps.
CONFIGURATION mode

```
snmp-server trap-source
```

Example of RFC-Defined SNMP Traps and Related Enable Commands

The following example lists the RFC-defined SNMP traps and the command used to enable each. The `coldStart` and `warmStart` traps are enabled using a single command.

```
snmp authentication      SNMP_AUTH_FAIL:SNMP Authentication failed.Request with
invalid community string.

snmp coldstart           SNMP_COLD_START: Agent Initialized - SNMP COLD_START.
                        SNMP_WARM_START:Agent Initialized - SNMP WARM_START.

snmp linkdown            PORT_LINKDN:changed interface state to down:%d


snmp linkup              PORT_LINKUP:changed interface state to up:%d
```

Enabling a Subset of SNMP Traps

You can enable a subset of Dell Networking enterprise-specific SNMP traps using one of the following listed command options.

To enable a subset of Dell Networking enterprise-specific SNMP traps, use the following command.

- Enable a subset of SNMP traps.
`snmp-server enable traps`

 **NOTE:** The `envmon` option enables all environment traps including those traps that are enabled with the `envmon supply`, `envmon temperature`, and `envmon fan` options.

Example of Dell Networking Enterprise-specific SNMP Traps

envmon temperature

```
MINOR_TEMP: Minor alarm: chassis temperature
MINOR_TEMP_CLR: Minor alarm cleared: chassis temperature normal (%s %d
temperature is within threshold of %dC)
MAJOR_TEMP: Major alarm: chassis temperature high (%s temperature reaches or
exceeds threshold of %dC)
MAJOR_TEMP_CLR: Major alarm cleared: chassis temperature lower (%s %d
temperature is within threshold of %dC)
```

xstp

```
%SPANMGR-5-STP_NEW_ROOT: New Spanning Tree Root, Bridge ID Priority 32768,
Address 0001.e801.fc35.
%SPANMGR-5-STP_TOPOLOGY_CHANGE: Bridge port GigabitEthernet 11/38 transitioned
from Forwarding to Blocking state.
%SPANMGR-5-MSTP_NEW_ROOT_BRIDGE: Elected root bridge for instance 0.
```

entity

Enable entity change traps

```
Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (1487406) 4:07:54.06,
SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::mib-2.47.2.0.1,
SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 4
Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (1488564) 4:08:05.64,
SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::mib-2.47.2.0.1,
SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 5
Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (1489064) 4:08:10.64,
SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::mib-2.47.2.0.1,
SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 6
Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (1489568)
4:08:15.68,SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::mib-2.47.2.0.1,
SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 7
```

<cr>

SNMP Copy Config Command Completed

```
%RPM0-P:CP %SNMP-4-RMON_RISING_THRESHOLD: RMON rising threshold alarm from SNMP
OID <oid>
%RPM0-P:CP %SNMP-4-RMON_FALLING_THRESHOLD: RMON falling threshold alarm from
SNMP OID <oid>
%RPM0-P:CP %SNMP-4-RMON_HC_RISING_THRESHOLD: RMON high-capacity rising threshold
alarm from SNMP OID <oid>
```

coldstart

warmstart linkdown

linkup

```
10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (6796)
0:01:07.96, SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-MIB::coldStart,
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.1 = INTEGER: 6,
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "SNMP_COLD_START: Agent
Initialized - SNMP COLD_START.", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 =
INTEGER: 1
10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks: (6756)
0:01:07.56, SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-MIB::warmStart,
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.1 = INTEGER: 6,
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "SNMP_WARM_START: Agent
Initialized - SNMP WARM_START.", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 =
INTEGER: 1
10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks:
(625882) 1:44:18.82, SNMPv2-MIB::snmpTrapOID.0 = OID: IF-MIB::linkUp,
IF-MIB::ifIndex.45158657 = INTEGER: 45158657,
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "OSTATE_UP: Changed interface
state to up: Te 0/43", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 14
10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks:
(645746) 1:47:37.46, SNMPv2-MIB::snmpTrapOID.0 = OID: IF-MIB::linkDown,
IF-MIB::ifIndex.45420801 = INTEGER: 45420801,
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "OSTATE_DN: Changed interface
state to down: Te 0/44", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 22
```

ets

ETS peer state enabled

```
10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks:
(625916) 1:44:19.16, SNMPv2-MIB::snmpTrapOID.0 = OID:
SNMPv2-SMI::enterprises.6027.3.15.4.0.3,
SNMPv2-SMI::enterprises.6027.3.15.4.1.1.0 = INTEGER: 45158657,
SNMPv2-SMI::enterprises.6027.3.15.4.1.2.0 = INTEGER: 1,
SNMPv2-SMI::enterprises.6027.3.15.4.0 = STRING:
"ETS_TRAP_TYPE_PEER_STATE_CHANGE: ETS Peer state changed to enabled for port Te
0/43", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 17
```

ETS peer state disabled

```
10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks:
(645772) 1:47:37.72, SNMPv2-MIB::snmpTrapOID.0 = OID:
```



```
SNMPv2-SMI::enterprises.6027.3.15.4.0.3,  
SNMPv2-SMI::enterprises.6027.3.15.4.1.1.0 = INTEGER: 45420801,  
SNMPv2-SMI::enterprises.6027.3.15.4.1.2.0 = INTEGER: 2,  
SNMPv2-SMI::enterprises.6027.3.15.4.0 = STRING:  
"ETS_TRAP_TYPE_PEER_STATE_CHANGE: ETS Peer state changed to disabled for port Te  
0/44", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 23
```

pfc


```
pfc peer state enabled  
10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks:  
(626100) 1:44:21.00, SNMPv2-MIB::snmpTrapOID.0 = OID:  
SNMPv2-SMI::enterprises.6027.3.15.4.0.7,  
SNMPv2-SMI::enterprises.6027.3.15.4.1.1.0 = INTEGER: 45420801,  
SNMPv2-SMI::enterprises.6027.3.15.4.1.2.0 = INTEGER: 1,  
SNMPv2-SMI::enterprises.6027.3.15.4.0 = STRING:  
"PFC_TRAP_TYPE_PEER_STATE_CHANGE: PFC Peer state changed to enabled for port Te  
0/44", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 21  
pfc peer state disabled  
10.16.130.140 [10.16.130.140]: Trap SNMPv2-MIB::sysUpTime.0 = Timeticks:  
(645794) 1:47:37.94, SNMPv2-MIB::snmpTrapOID.0 = OID:  
SNMPv2-SMI::enterprises.6027.3.15.4.0.7,  
SNMPv2-SMI::enterprises.6027.3.15.4.1.1.0 = INTEGER: 45420801,  
SNMPv2-SMI::enterprises.6027.3.15.4.1.2.0 = INTEGER: 2,  
SNMPv2-SMI::enterprises.6027.3.15.4.0 = STRING:  
"PFC_TRAP_TYPE_PEER_STATE_CHANGE: PFC Peer state changed to disabled for port Te  
0/44", SNMPv2-SMI::enterprises.6027.3.6.1.1.2.0 = INTEGER: 24
```

Enabling an SNMP Agent to Notify Syslog Server Failure

You can configure a network device to send an SNMP trap in the event of an audit processing failure due to connectivity issues with the syslog server.

If a connectivity failure occurs on a syslog server that is configured for reliable transmission, an SNMP trap is sent. Also a message is written to the console indicating that the configured syslog server has failed.

The SNMP trap is sent only periodically. Meaning, when a syslog connection fails and the time-interval between the last syslog notification and current time is greater than or equal to five minutes, a trap is sent. This restriction applies to the console message also.

 **NOTE:** If a syslog server failure event is generated before the SNMP agent service starts, then SNMP trap is not sent successfully.

To enable an SNMP agent to send a trap when the syslog server is not reachable, use the following command:

```
CONFIGURATION MODE
```

```
snmp-server enable traps snmp syslog-unreachable
```

To enable an SNMP agent to send a trap when the syslog server resumes connectivity, use the following command:

```
CONFIGURATION MODE
```

```
snmp-server enable traps snmp syslog-reachable
```

Table 59. List of Syslog Server MIBs that have read access

MIB Object	OID	Object Values	Description
dF10SysLogTraps	1.3.6.1.4.1.6027.3.30.1.1	1 = reachable 2 = unreachable	Specifies whether the syslog server is reachable or unreachable.

Following example shows the SNMP trap that is sent when connectivity to the syslog server is lost:

```
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (19738) 0:03:17.38      SNMPv2-
MIB::snmpTrapOID.0 = OID: SNMPv2-
SMI::enterprises.6027.3.30.1.1.1      SNMPv2-SMI::enterprises.6027.3.30.1.1 =
STRING: "NOT_REACHABLE: Syslog server
10.11.226.121 (port: 9140) is not reachable" SNMPv2-SMI::enterprises.
6027.3.6.1.1.2.0 = INTEGER: 2
```

Following is the sample auditlog message that is received by other syslog servers, which are reachable:

```
Oct 21 00:46:13: dv-fedgov-s4810-6: %EVL-6-NOT_REACHABLE:Syslog server
10.11.226.121 (port: 9140) is not reachable
```

Following example shows the SNMP trap that is sent when connectivity to the syslog server is resumed:

```
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (10230) 0:01:42.30      SNMPv2-
MIB::snmpTrapOID.0 = OID: SNMPv2-
SMI::enterprises.6027.3.30.1.1.2      SNMPv2-SMI::enterprises.6027.3.30.1.1 =
STRING: "REACHABLE: Syslog server
10.11.226.121 (port: 9140) is reachable"SNMPv2-SMI::enterprises.
6027.3.6.1.1.2.0 = INTEGER: 2
```

Following is the sample auditlog message that is received by other syslog servers, which are reachable:

```
Oct 21 05:26:04: dv-fedgov-s4810-6: %EVL-6-REACHABLE:Syslog server
10.11.226.121 (port: 9140) is reachable
```

Copy Configuration Files Using SNMP

To do the following, use SNMP from a remote client.

- copy the running-config file to the startup-config file
- copy configuration files from the Dell Networking system to a server
- copy configuration files from a server to the Dell Networking system

You can perform all of these tasks using IPv4 addresses.

The following table lists the relevant MIBs for these functions are.

Table 60. MIB Objects for Copying Configuration Files via SNMP

MIB Object	OID	Object Values	Description
copySrcFileType	. 1.3.6.1.4.1.6027.3.5.1.1.1. 1.2	1 = Dell Networking OS file 2 = running-config	Specifies the type of file to copy from. The range is: <ul style="list-style-type: none"> • If copySrcFileType is running-config or

MIB Object	OID	Object Values	Description
		3 = startup-config	startup-config, the default copySrcFileLocation is flash. <ul style="list-style-type: none"> If copySrcFileType is a binary file, you must also specify copySrcFileLocation and copySrcFileName.
copySrcFileLocation	. 1.3.6.1.4.1.6027.3.5.1.1.1.1.3	1 = flash 2 = n/a 3 = tftp 4 = ftp 5 = scp 6 = usbflash	Specifies the location of source file. <ul style="list-style-type: none"> If copySrcFileLocation is FTP or SCP, you must specify copyServerAddress, copyUserName, and copyUserPassword.
copySrcFileName	. 1.3.6.1.4.1.6027.3.5.1.1.1.1.4	Path (if the file is not in the current directory) and filename.	Specifies name of the file. <ul style="list-style-type: none"> If copySourceFileType is set to running-config or startup-config, copySrcFileName is not required.
copyDestFileType	. 1.3.6.1.4.1.6027.3.5.1.1.1.1.5	1 = Dell Networking OS file 2 = running-config 3 = startup-config	Specifies the type of file to copy to. <ul style="list-style-type: none"> If copySourceFileType is running-config or startup-config, the default copyDestFileLocation is flash. If copyDestFileType is a binary, you must specify copyDestFileLocation and copyDestFileName.
copyDestFileLocation	. 1.3.6.1.4.1.6027.3.5.1.1.1.1.6	1 = flash 2 = n/a 3 = tftp 4 = ftp	Specifies the location of destination file. <ul style="list-style-type: none"> If copyDestFileLocation is FTP or SCP, you must specify copyServerAddress,

MIB Object	OID	Object Values	Description
		5 = scp	copyUserName, and copyUserPassword.
copyDestFileName	. 1.3.6.1.4.1.6027.3.5.1.1.1. 1.7	Path (if the file is not in the default directory) and filename.	Specifies the name of destination file.
copyServerAddress	. 1.3.6.1.4.1.6027.3.5.1.1.1. 1.8	IP Address of the server.	The IP address of the server. <ul style="list-style-type: none"> If you specify copyServerAddress, you must also specify copyUserName and copyUserPassword.
copyUserName	. 1.3.6.1.4.1.6027.3.5.1.1.1. 1.9	Username for the server.	Username for the FTP, TFTP, or SCP server. <ul style="list-style-type: none"> If you specify copyUserName, you must also specify copyUserPassword.
copyUserPassword	. 1.3.6.1.4.1.6027.3.5.1.1.1. 1.10	Password for the server.	Password for the FTP, TFTP, or SCP server.

Copying a Configuration File

To copy a configuration file, use the following commands.

1. Create an SNMP community string with read/write privileges.

CONFIGURATION mode

```
snmp-server community community-name rw
```


2. Copy the *f10-copy-config.mib* MIB from the Dell iSupport web page to the server to which you are copying the configuration file.
3. On the server, use the `snmpset` command as shown in the following example.

```
snmpset -v snmp-version -c community-name -m mib_path/f10-copy-config.mib force10system-ip-address mib-object.index {i | a | s} object-value...
```

- Every specified object must have an object value and must precede with the keyword `i`. Refer to the previous table.
- *index* must be unique to all previously executed `snmpset` commands. If an index value has been used previously, a message like the following appears. In this case, increment the index value and enter the command again.

- To complete the command, use as many MIB Objects in the command as required by the MIB Object descriptions.

```
Error in packet.
Reason: notWritable (that object does not support modification)
Failed object: FTOS-COPY-CONFIG-MIB::copySrcFileType.101
```

 **NOTE:** You can use the entire OID rather than the object name. Use the form: *OID.index i object-value*.

To view more information, use the following options in the `snmpset` command.

- `-c`: View the community, either public or private.
- `-m`: View the MIB files for the SNMP command.
- `-r`: Number of retries using the option
- `-t`: View the timeout.
- `-v`: View the SNMP version (either 1, 2, 2d, or 3).

The following examples show the `snmpset` command to copy a configuration. These examples assume that:

- the server OS is UNIX
- you are using SNMP version 2c
- the community name is public
- the file `f10-copy-config.mib` is in the current directory or in the `snmpset` tool path

Copying Configuration Files via SNMP

To copy the running-config to the startup-config from the UNIX machine, use the following command.

- Copy the running-config to the startup-config from the UNIX machine.


```
snmpset -v 2c -c public -m ./f10-copy-config.mif force10system-ip-address
copySrcFileType.index i 2 copyDestFileType.index i 3
```

Example of Copying Configuration Files (Using MIB Object Names)

Example of Copying Configuration Files (Using OIDs)

The following examples show the command syntax using MIB object names and the same command using the object OIDs. In both cases, a unique index number follows the object.

```
> snmpset -v 2c -r 0 -t 60 -c public -m ./f10-copy-config.mib 10.10.10.10
copySrcFileType.101 i
2 copyDestFileType.101 i 3
FORCE10-COPY-CONFIG-MIB::copySrcFileType.101 = INTEGER: runningConfig(2)
FORCE10-COPY-CONFIG-MIB::copyDestFileType.101 = INTEGER: startupConfig(3)
```

```
> snmpset -v 2c -c public -m ./f10-copy-config.mib 10.10.10.10
.1.3.6.1.4.1.6027.3.5.1.1.1.2.100 i 2 .1.3.6.1.4.1.6027.3.5.1.1.1.5.100 i 3
FORCE10-COPY-CONFIG-MIB::copySrcFileType.100 = INTEGER: runningConfig(2)
FORCE10-COPY-CONFIG-MIB::copyDestFileType.100 = INTEGER: startupConfig(3)
```

Copying the Startup-Config Files to the Running-Config

To copy the startup-config to the running-config from a UNIX machine, use the following command.

- Copy the startup-config to the running-config from a UNIX machine.

```
snmpset -c private -v 2c force10system-ip-address copySrcFileType.index i 3  
copyDestFileType.index i 2
```

Example of Copying Configuration Files from a UNIX Machine (Using Object Name)

Example of Copying Configuration Files from a UNIX Machine (Using OID)

```
> snmpset -c public -v 2c -m ./f10-copy-config.mib 10.11.131.162  
copySrcFileType.7 i 3  
copyDestFileType.7 i 2  
FORCE10-COPY-CONFIG-MIB::copySrcFileType.7 = INTEGER: runningConfig(3)  
FORCE10-COPY-CONFIG-MIB::copyDestFileType.7 = INTEGER: startupConfig(2)  
  
>snmpset -c public -v 2c 10.11.131.162 .1.3.6.1.4.1.6027.3.5.1.1.1.2.8 i 3  
.1.3.6.1.4.1.6027.3.5.1.1.1.5.8 i 2  
SNMPv2-SMI::enterprises.6027.3.5.1.1.1.2.8 = INTEGER: 3  
SNMPv2-SMI::enterprises.6027.3.5.1.1.1.5.8 = INTEGER: 2
```

Copying the Startup-Config Files to the Server via FTP

To copy the startup-config to the server via FTP from the UNIX machine, use the following command.

Copy the startup-config to the server via FTP from the UNIX machine.

```
snmpset -v 2c -c public -m ./f10-copy-config.mib force10system-ip-address  
copySrcFileType.index i 2 copyDestFileName.index s filepath/filename  
copyDestFileLocation.index i 4 copyServerAddress.index a server-ip-address  
copyUserName.index s server-login-id copyUserPassword.index s server-login-  
password
```


- precede *server-ip-address* by the keyword *a*.
- precede the values for *copyUsername* and *copyUserPassword* by the keyword *s*.

Example of Copying Configuration Files via FTP From a UNIX Machine

```
> snmpset -v 2c -c private -m ./f10-copy-config.mib 10.10.10.10 copySrcFileType.  
110 i 2  
copyDestFileName.110 s /home/startup-config copyDestFileLocation.110 i 4  
copyServerAddress.110  
a 11.11.11.11 copyUserName.110 s mylogin copyUserPassword.110 s mypass  
FORCE10-COPY-CONFIG-MIB::copySrcFileType.110 = INTEGER: runningConfig(2)  
FORCE10-COPY-CONFIG-MIB::copyDestFileName.110 = STRING: /home/startup-config  
FORCE10-COPY-CONFIG-MIB::copyDestFileLocation.110 = INTEGER: ftp(4)  
FORCE10-COPY-CONFIG-MIB::copyServerAddress.110 = IPAddress: 11.11.11.11  
FORCE10-COPY-CONFIG-MIB::copyUserName.110 = STRING: mylogin  
FORCE10-COPY-CONFIG-MIB::copyUserPassword.110 = STRING: mypass
```

Copying the Startup-Config Files to the Server via TFTP

To copy the startup-config to the server via TFTP from the UNIX machine, use the following command.

 **NOTE:** Verify that the file exists and its permissions are set to 777. Specify the relative path to the TFTP root directory.

- Copy the startup-config to the server via TFTP from the UNIX machine.

```
snmpset -v 2c -c public -m ./f10-copy-config.mib force10system-ip-address
copySrcFileType.index i 3 copyDestFileType.index i 1 copyDestFileName.index s
filepath/filename copyDestFileLocation.index i 3 copyServerAddress.index a
server-ip-address
```

Example of Copying Configuration Files via TFTP From a UNIX Machine

```
.snmpset -v 2c -c private -m ./f10-copy-config.mib 10.10.10.10
copySrcFileType.4 i 3
copyDestFileType.4 i 1
copyDestFileLocation.4 i 3
copyDestFileName.4 s /home/myfilename
copyServerAddress.4 a 11.11.11.11
```

Copying a Binary File to the Startup-Configuration

To copy a binary file from the server to the startup-configuration on the Dell Networking system via FTP, use the following command.

- Copy a binary file from the server to the startup-configuration on the Dell Networking system via FTP.

```
snmpset -v 2c -c public -m ./f10-copy-config.mib force10system-ip-address
copySrcFileType.index i 1 copySrcFileLocation.index i 4 copySrcFileName.index
s filepath/filename copyDestFileType.index i 3 copyServerAddress.index a
server-ip-address copyUserName.index s server-login-id copyUserPassword.index
s server-login-password
```

Example of Copying a Binary File From the Server to the Startup-Configuration via FTP

```
> snmpset -v 2c -c private -m ./f10-copy-config.mib 10.10.10.10 copySrcFileType.
10 i 1
copySrcFileLocation.10 i 4 copyDestFileType.10 i 3 copySrcFileName.10 s /home/
myfilename
copyServerAddress.10 a 172.16.1.56 copyUserName.10 s mylogin copyUserPassword.
10 s mypass
```

Additional MIB Objects to View Copy Statistics

Dell Networking provides more MIB objects to view copy statistics, as shown in the following table.

Table 61. Additional MIB Objects for Copying Configuration Files via SNMP

MIB Object	OID	Values	Description
copyState	. 1.3.6.1.4.1.6027.3.5.1.1.1. 1.11	1= running 2 = successful 3 = failed	Specifies the state of the copy operation.
copyTimeStarted	. 1.3.6.1.4.1.6027.3.5.1.1.1. 1.12	Time value	Specifies the point in the up-time clock that the copy operation started.
copyTimeCompleted	. 1.3.6.1.4.1.6027.3.5.1.1.1. 1.13	Time value	Specifies the point in the up-time clock that the copy operation completed.
copyFailCause	. 1.3.6.1.4.1.6027.3.5.1.1.1. 1.14	1 = bad filename 2 = copy in progress 3 = disk full 4 = file exists 5 = file not found 6 = timeout 7 = unknown	Specifies the reason the copy request failed.
copyEntryRowStatus	. 1.3.6.1.4.1.6027.3.5.1.1.1. 1.15	Row status	Specifies the state of the copy operation. Uses CreateAndGo when you are performing the copy. The state is set to <i>active</i> when the copy is completed.

MIB Support to Display the Available Memory Size on Flash

Dell Networking provides more MIB objects to display the available memory size on flash memory. The following table lists the MIB object that contains the available memory size on flash memory.

Table 62. MIB Objects for Displaying the Available Memory Size on Flash via SNMP

MIB Object	OID	Description
chStackUnitFlashUsageUtil	1.3.6.1.4.1.6027.3.19.1.2.8.1.6	Contains flash memory usage in percentage.

The chStackUnitUtilTable MIB table contains the chStackUnitFlashUsageUtil MIB object which contains the flash memory usage percent.

Viewing the Available Flash Memory Size

- To view the available flash memory using SNMP, use the following command.

```
snmpget -v2c -c public 192.168.60.120 .1.3.6.1.4.1.6027.3.10.1.2.9.1.6.1  
enterprises.6027.3.10.1.2.9.1.5.1 = Gauge32: 24
```

The output above displays that 24% of the flash memory is used.

MIB Support to Display the Software Core Files Generated by the System

Dell Networking provides MIB objects to display the software core files generated by the system. The chSysSwCoresTable contains the list of software core files generated by the system. The following table lists the related MIB objects.

Table 63. MIB Objects for Displaying the Software Core Files Generated by the System

MIB Object	OID	Description
chSysSwCoresTable	1.3.6.1.4.1.6027.3.19.1.2.9	This is the table that contains the list of software core files generated by the system.
chSysCoresEntry	1.3.6.1.4.1.6027.3.19.1.2.9.1	Entry number.
chSysCoresInstance	1.3.6.1.4.1.6027.3.19.1.2.9.1.1	Stores the indexed information about the available software core files.
chSysCoresFileName	1.3.6.1.4.1.6027.3.19.1.2.9.1.2	Contains the core file names and the file paths.

MIB Object	OID	Description
chSysCoresTimeCreated	1.3.6.1.4.1.6027.3.19.1.2.9.1.3	Contains the time at which core files are created.
chSysCoresStackUnitNumber	1.3.6.1.4.1.6027.3.19.1.2.9.1.4	Contains information that includes which stack unit or processor the core file was originated from.
chSysCoresProcess	1.3.6.1.4.1.6027.3.19.1.2.9.1.5	Contains information that includes the process names that generated each core file.

Viewing the Software Core Files Generated by the System

- To view the viewing the software core files generated by the system, use the following command.

```
snmpwalk -v2c -c public 192.168.60.120 .1.3.6.1.4.1.6027.3.10.1.2.10

enterprises.6027.3.10.1.2.10.1.1.1.1 = 1
enterprises.6027.3.10.1.2.10.1.1.1.2 = 2
enterprises.6027.3.10.1.2.10.1.1.1.3 = 3
enterprises.6027.3.10.1.2.10.1.1.2.1 = 1
enterprises.6027.3.10.1.2.10.1.2.1.1 = "/CORE_DUMP_DIR/flashmnr.core.gz"
enterprises.6027.3.10.1.2.10.1.2.1.2 = "/CORE_DUMP_DIR/FTP_STK_MEMBER/
f10cp_l2mgr_131108080758_Stk1.acore.gz"
enterprises.6027.3.10.1.2.10.1.2.1.3 = "/CORE_DUMP_DIR/FTP_STK_MEMBER/
f10cp_vrrp_140522124357_Stk1.acore.gz"
enterprises.6027.3.10.1.2.10.1.2.2.1 =
"/CORE_DUMP_DIR/FTP_STK_MEMBER/f10cp_sysd_140617134445_Stk0.acore.gz"
enterprises.6027.3.10.1.2.10.1.3.1.1 = "Fri Mar 14 11:51:46 2014"
enterprises.6027.3.10.1.2.10.1.3.1.2 = "Fri Nov 8 08:11:16 2013"
enterprises.6027.3.10.1.2.10.1.3.1.3 = "Fri May 23 05:05:16 2014"
enterprises.6027.3.10.1.2.10.1.3.2.1 = "Tue Jun 17 14:19:26 2014"
enterprises.6027.3.10.1.2.10.1.4.1.1 = 0
enterprises.6027.3.10.1.2.10.1.4.1.2 = 1
enterprises.6027.3.10.1.2.10.1.4.1.3 = 1
enterprises.6027.3.10.1.2.10.1.4.2.1 = 0
enterprises.6027.3.10.1.2.10.1.5.1.1 = "flashmnr"
enterprises.6027.3.10.1.2.10.1.5.1.2 = "l2mgr"
enterprises.6027.3.10.1.2.10.1.5.1.3 = "vrrp" Hex: 76 72 72 70
enterprises.6027.3.10.1.2.10.1.5.2.1 = "sysd" Hex: 73 79 73 64
```

The output above displays that the software core files generated by the system.


Obtaining a Value for MIB Objects

To obtain a value for any of the MIB objects, use the following command.

- Get a copy-config MIB object value.

```
snmpset -v 2c -c public -m /f10-copy-config.mib force10system-ip-address
[OID.index | mib-object.index]
```

index: the index value used in the `snmpset` command used to complete the copy operation.


 **NOTE:** You can use the entire OID rather than the object name. Use the form: *OID.index*.

Example of Getting a MIB Object Value (Using Object Name)

Example of Getting a MIB Object Value (Using OID)

The following examples show the `snmpget` command to obtain a MIB object value. These examples assume that:

- the server OS is UNIX
- you are using SNMP version 2c
- the community name is public
- the file `f10-copy-config.mib` is in the current directory

 **NOTE:** In UNIX, enter the `snmpset` command for help using this command.

The following examples show the command syntax using MIB object names and the same command using the object OIDs. In both cases, the same index number used in the `snmpset` command follows the object.

```
> snmpget -v 2c -c private -m ./f10-copy-config.mib 10.11.131.140
copyTimeCompleted.110
FORCE10-COPY-CONFIG-MIB::copyTimeCompleted.110 = Timeticks: (1179831) 3:16:38.31
```

```
> snmpget -v 2c -c private 10.11.131.140 .1.3.6.1.4.1.6027.3.5.1.1.1.1.13.110
SNMPv2-SMI::enterprises.6027.3.5.1.1.1.1.13.110 = Timeticks: (1179831)
3:16:38.31
```

Manage VLANs using SNMP

The `qBridgeMIB` managed objects in `Q-BRIDGE-MIB`, defined in RFC 2674, allows you to use SNMP to manage VLANs.

Creating a VLAN

To create a VLAN, use the `dot1qVlanStaticRowStatus` object.

The `snmpset` operation shown in the following example creates VLAN 10 by specifying a value of 4 for instance 10 of the `dot1qVlanStaticRowStatus` object.

Example of Creating a VLAN using SNMP

```
> snmpset -v2c -c mycommunity 123.45.6.78 .1.3.6.1.2.1.17.7.1.4.3.1.5.10 i 4
SNMPv2-SMI::mib-2.17.7.1.4.3.1.5.10 = INTEGER: 4
```

Assigning a VLAN Alias

Write a character string to the `dot1qVlanStaticName` object to assign a name to a VLAN.

Example of Assigning a VLAN Alias using SNMP

[Unix system output]

```
> snmpset -v2c -c mycommunity 10.11.131.185 .
1.3.6.1.2.1.17.7.1.4.3.1.1.1107787786 s "My
VLAN"
SNMPv2-SMI::mib-2.17.7.1.4.3.1.1.1107787786 = STRING: "My VLAN"
```

[Dell Networking system output]

```
Dell#show int vlan 10
Vlan 10 is down, line protocol is down
Vlan alias name is: My VLAN
Address is 00:01:e8:cc:cc:ce, Current address is 00:01:e8:cc:cc:ce
Interface index is 1107787786
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed auto
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 01:01:00
Queueing strategy: fifo
Time since last interface status change: 01:01:00
```

Displaying the Ports in a VLAN

Dell Networking OS identifies VLAN interfaces using an interface index number that is displayed in the output of the `show interface vlan` command.

Example of Viewing the VLAN Interface Index Number Using SNMP

Example of Viewing VLAN Ports Using SNMP (No Ports Assigned)

Example of Viewing VLAN Ports Using SNMP (Port Assigned)

```
Dell(conf)#do show interface vlan 10
Vlan 10 is down, line protocol is down
Address is 00:01:e8:cc:cc:ce, Current address is 00:01:e8:cc:cc:ce
Interface index is 1107787786
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed auto
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:12:42
Queueing strategy: fifo
Time since last interface status change: 00:12:42
```

To display the ports in a VLAN, send an `snmpget` request for the object `dot1qStaticEgressPorts` using the interface index as the instance number, as shown for an S-Series.

```
> snmpget -v2c -c mycommunity 10.11.131.185 .
1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786
SNMPv2-SMI::mib-2.17.7.1.4.3.1.2.1107787786 = Hex-STRING:
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

The table that the Dell Networking system sends in response to the `snmpget` request is a table that contains hexadecimal (hex) pairs, each pair representing a group of eight ports.

- Seven hex pairs represent a stack unit. Seven pairs accommodate the greatest number of ports available — 64 ports on the MXL switch, the last stack unit is assigned eight pairs, the eight pair is unused.

The first hex pair, 00 in the previous example, represents ports 1 to 7 in Stack Unit 0. The next pair to the right represents ports 8 to 15. To resolve the hex pair into a representation of the individual ports, convert the hex pair to binary. Consider the first hex pair 00, which resolves to 0000 0000 in binary:

- Each position in the 8-character string is for one port, starting with Port 1 at the left end of the string, and ending with Port 8 at the right end. A 0 indicates that the port is not a member of the VLAN; a 1 indicates VLAN membership.

All hex pairs are 00, indicating that no ports are assigned to VLAN 10. In the following example, Port 0/2 is added to VLAN 10 as untagged; the first hex pair changes from 00 to 04.

[Dell Networking system output]

```
R5(conf)#do show vlan id 10
```


```
Codes: * - Default VLAN, G - GVRP VLANs
Q: U - Untagged, T - Tagged
    x - Dot1x untagged, X - Dot1x tagged
    G - GVRP tagged, M - Vlan-stack
```

NUM	Status	Description	Q	Ports
10	Inactive		U	Gi 0/2

[Unix system output]

```
> snmpget -v2c -c mycommunity 10.11.131.185 .
1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786
SNMPv2-SMI::mib-2.17.7.1.4.3.1.2.1107787786 = Hex-STRING: 40 00 00 00 00 00 00 00
00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
```

The value 40 is in the first set of 7 hex pairs, indicating that these ports are in Stack Unit 0. The hex value 40 is 0100 0000 in binary. As described, the left-most position in the string represents Port 1. The next position from the left represents Port 2 and has a value of 1, indicating that Port 0/2 is in VLAN 10. The remaining positions are 0, so those ports are not in the VLAN.


 **NOTE:** The table contains none of the other information the command provides, such as port speed or whether the ports are tagged or untagged.

Add Tagged and Untagged Ports to a VLAN

The value `dot1qVlanStaticEgressPorts` object is an array of all VLAN members.

The `dot1qVlanStaticUntaggedPorts` object is an array of only untagged VLAN members. All VLAN members that are not in `dot1qVlanStaticUntaggedPorts` are tagged.

- To add a tagged port to a VLAN, write the port to the `dot1qVlanStaticEgressPorts` object.
- To add an untagged port to a VLAN, write the port to the `dot1qVlanStaticEgressPorts` and `dot1qVlanStaticUntaggedPorts` objects.

 **NOTE:** Whether adding a tagged or untagged port, specify values for both `dot1qVlanStaticEgressPorts` and `dot1qVlanStaticUntaggedPorts`.

In the following example, Port 0/2 is added as an untagged member of VLAN 10.

Example of Adding an Untagged Port to a VLAN using SNMP

Example of Adding a Tagged Port to a VLAN using SNMP

```
>snmpset -v2c -c mycommunity 10.11.131.185 .
1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786 x "40 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00"
.1.3.6.1.2.1.17.7.1.4.3.1.4.1107787786 x "40 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00"
SNMPv2-SMI::mib-2.17.7.1.4.3.1.2.1107787786 = Hex-STRING: 40 00 00 00 00 00 00
00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
SNMPv2-SMI::mib-2.17.7.1.4.3.1.4.1107787786 = Hex-STRING: 40 00 00 00 00 00 00
00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

In the following example, Port 0/2 is added as a tagged member of VLAN 10.

```
>snmpset -v2c -c mycommunity 10.11.131.185 .
1.3.6.1.2.1.17.7.1.4.3.1.2.1107787786 x "40 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00"
.1.3.6.1.2.1.17.7.1.4.3.1.4.1107787786 x "00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00"
SNMPv2-SMI::mib-2.17.7.1.4.3.1.2.1107787786 = Hex-STRING: 40 00 00 00 00 00 00
00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00
SNMPv2-SMI::mib-2.17.7.1.4.3.1.4.1107787786 = Hex-STRING: 00 00 00 00 00 00 00
00 00 00 00
00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

Enabling and Disabling a Port using SNMP

To enable and disable a port using SNMP, use the following commands.

1. Create an SNMP community on the Dell system.
CONFIGURATION mode

```
snmp-server community
```

- From the Dell Networking system, identify the interface index of the port for which you want to change the admin status.

EXEC Privilege mode

```
show interface
```

Or, from the management system, use the `snmpwalk` command to identify the interface index.

- Enter the `snmpset` command to change the admin status using either the object descriptor or the OID.

```
snmpset with descriptor: snmpset -v version -c community agent-ip  
ifAdminStatus.ifindex i {1 | 2}
```

```
snmpset with OID: snmpset -v version -c community agent-ip .  
1.3.6.1.2.1.2.2.1.7.ifindex i {1 | 2}
```

Choose integer 1 to change the admin status to Up, or 2 to change the admin status to Down.

Fetch Dynamic MAC Entries using SNMP

Dell Networking supports the RFC 1493 dot1d table for the default VLAN and the dot1q table for all other VLANs.


 **NOTE:** The 802.1q Q-BRIDGE MIB defines VLANs regarding 802.1d, as 802.1d itself does not define them. As a switchport must belong a VLAN (the default VLAN or a configured VLAN), all MAC address learned on a switchport are associated with a VLAN. For this reason, the Q-Bridge MIB is used for MAC address query. Moreover, specific to MAC address query, the MAC address indexes dot1dTpFdbTable only for a single forwarding database, while dot1qTpFdbTable has two indices — VLAN ID and MAC address — to allow for multiple forwarding databases and considering that the same MAC address is learned on multiple VLANs. The VLAN ID is added as the first index so that MAC addresses are read by the VLAN, sorted lexicographically. The MAC address is part of the OID instance, so in this case, lexicographic order is according to the most significant octet.

Table 64. MIB Objects for Fetching Dynamic MAC Entries in the Forwarding Database

MIB Object	OID	MIB	Description
dot1dTpFdbTable	.1.3.6.1.2.1.17.4.3	Q-BRIDGE MIB	List the learned unicast MAC addresses on the default VLAN.
dot1qTpFdbTable	.1.3.6.1.2.1.17.7.1.2. 2	Q-BRIDGE MIB	List the learned unicast MAC addresses on non-default VLANs.
dot3aCurAggFdb Table	.1.3.6.1.4.1.6027.3.2. 1.1.5	F10-LINK-AGGREGATION -MIB	List the learned MAC addresses of aggregated links (LAG).

In the following example, R1 has one dynamic MAC address, learned off of port GigabitEthernet 1/21, which a member of the default VLAN, VLAN 1. The SNMP walk returns the values for dot1dTpFdbAddress, dot1dTpFdbPort, and dot1dTpFdbStatus.

Each object is comprised of an OID concatenated with an instance number. In the case of these objects, the instance number is the decimal equivalent of the MAC address; derive the instance number by converting each hex pair to its decimal equivalent. For example, the decimal equivalent of E8 is 232, and so the instance number for MAC address 00:01:e8:06:95:ac is.0.1.232.6.149.172.

The value of dot1dTpFdbPort is the port number of the port off which the system learns the MAC address. In this case, of TenGigabitEthernet 1/21, the manager returns the integer 118.

Example of Fetching MAC Addresses Learned on the Default VLAN Using SNMP

Example of Fetching MAC Addresses Learned on a Non-default VLAN Using SNMP

Example of Fetching MAC Addresses Learned on a Port-Channel Using SNMP

```
-----MAC Addresses on Dell Networking
System-----
Dell#show mac-address-table
VlanId Mac Address Type Interface State
1 00:01:e8:06:95:ac Dynamic Tengig 1/21 Active
-----Query from Management
Station-----
>snmpwalk -v 2c -c techpubs 10.11.131.162 .1.3.6.1.2.1.17.4.3.1
SNMPv2-SMI::mib-2.17.4.3.1.1.0.1.232.6.149.172 = Hex-STRING: 00 01 E8 06 95 AC
SNMPv2-SMI::mib-2.17.4.3.1.2.0.1.232.6.149.172 = INTEGER: 118
SNMPv2-SMI::mib-2.17.4.3.1.3.0.1.232.6.149.172 = INTEGER: 3
```

In the following example, GigabitEthernet 1/21 is moved to VLAN 1000, a non-default VLAN. To fetch the MAC addresses learned on non-default VLANs, use the object dot1qTpFdbTable. The instance number is the VLAN number concatenated with the decimal conversion of the MAC address.

```
-----MAC Addresses on Dell Networking
System-----
Dell#show mac-address-table
VlanId Mac Address Type Interface State
1000 00:01:e8:06:95:ac Dynamic Tengig 1/21 Active
-----Query from Management
Station-----
>snmpwalk -v 2c -c techpubs 10.11.131.162 .1.3.6.1.2.1.17.7.1.2.2.1
SNMPv2-SMI::mib-2.17.7.1.2.2.1.2.1000.0.1.232.6.149.172 = INTEGER: 118
SNMPv2-SMI::mib-2.17.7.1.2.2.1.3.1000.0.1.232.6.149.172 = INTEGER: 3
```

Use dot3aCurAggFdbTable to fetch the learned MAC address of a port-channel. The instance number is the decimal conversion of the MAC address concatenated with the port-channel number.

```
-----MAC Addresses on Dell Networking
System-----
Dell(conf)#do show mac-address-table
VlanId Mac Address Type Interface State
1000 00:01:e8:06:95:ac Dynamic Po 1 Active
-----Query from Management
Station-----
>snmpwalk -v 2c -c techpubs 10.11.131.162 .1.3.6.1.4.1.6027.3.2.1.1.5
SNMPv2-SMI::enterprises.6027.3.2.1.1.5.1.1.1000.0.1.232.6.149.172.1 = INTEGER:
1000
SNMPv2-SMI::enterprises.6027.3.2.1.1.5.1.2.1000.0.1.232.6.149.172.1 = Hex-
STRING: 00 01 E8
06 95 AC
```


SNMPv2-SMI::enterprises.6027.3.2.1.1.5.1.3.1000.0.1.232.6.149.172.1 = INTEGER: 1
SNMPv2-SMI::enterprises.6027.3.2.1.1.5.1.4.1000.0.1.232.6.149.172.1 = INTEGER: 1

Deriving Interface Indices


The Dell Networking OS assigns an interface number to each (configured or unconfigured) physical and logical interface.

The interface index is a binary number with bits that indicate the slot number, port number, interface type, and card type of the interface. The Dell Networking OS converts this binary index number to decimal, and displays it in the output of the `show interface` command.

Starting from the least significant bit (LSB):

- the first 14 bits represent the card type
- the next 4 bits represent the interface type
- the next 7 bits represent the port number
- the next 5 bits represent the slot number
- the next 1 bit is 0 for a physical interface and 1 for a logical interface
- the next 1 bit is unused

For example, the index 72925242 is 100010110001100000000111010 in binary. The binary interface index for GigabitEthernet 1/21 of a 48-port 10/100/1000Base-T line card with RJ-45 interface. Notice that the physical/logical bit and the final, unused bit are not given. The interface is physical, so represent this type of interface by a 0 bit, and the unused bit is always 0. These 2 bits are not given because they are the most significant bits, and leading zeros are often omitted.

 **NOTE:** The interface index does not change if the interface reloads or fails over. If the unit is renumbered (for any reason) the interface index changes during a reload.

To display the interface number, use the following command.

- Display the interface index number.
EXEC Privilege mode

```
show interface
```

Example of Deriving the Interface Index Number

To view the system image on Flash Partition A, use the `chSysSwInPartitionAlmgVers` object or, to view the system image on Flash Partition B, use the `chSysSwInPartitionBlmgVers` object.

Table 65. MIB Objects for Viewing the System Image on Flash Partitions

MIB Object	OID	Description	MIB
chSysSwInPartitionAlmgVers	1.3.6.1.4.1.6027.3.10.1.2.8.1.11	List the version string of the system image in Flash Partition A.	Chassis MIB
chSysSwInPartitionBlmgVers	1.3.6.1.4.1.6027.3.10.1.2.8.1.12	List the version string of the system image in Flash Partition B.	Chassis MIB

The system image can also be retrieved by performing an SNMP walk on the following OID: MIB Object is chSysSwModuleTable and the OID is 1.3.6.1.4.1.6027.3.10.1.2.8.

```
Dell#show interface tengig 1/21
TenGigabitEthernet 1/21 is up, line protocol is up
Hardware is Dell Force10Eth, address is 00:01:e8:0d:b7:4e
  Current address is 00:01:e8:0d:b7:4e
Interface index is 72925242
[output omitted]
```

Monitor Port-Channels

To check the status of a Layer 2 port-channel, use f10LinkAggMib (1.3.6.1.4.1.6027.3.2). In the following example, Po 1 is a switchport and Po 2 is in Layer 3 mode.

Example of SNMP Trap for Monitored Port-Channels

Example of Viewing Status of Learned MAC Addresses

Example of Viewing Changed Interface State for Monitored Ports

```
[senthilnathan@lithium ~]$ snmpwalk -v 2c -c public 10.11.1.1 .
1.3.6.1.4.1.6027.3.2.1.1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.1 = INTEGER: 1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.2 = INTEGER: 2
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.2.1 = Hex-STRING: 00 01 E8 13 A5 C7
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.2.2 = Hex-STRING: 00 01 E8 13 A5 C8
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.3.1 = INTEGER: 1107755009
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.3.2 = INTEGER: 1107755010
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.4.1 = INTEGER: 1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.4.2 = INTEGER: 1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.5.1 = Hex-STRING: 00 00
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.5.2 = Hex-STRING: 00 00
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.6.1 = STRING: "Gi 5/84 " << Channel member for Po1
SNMPv2-SMI::enterprises.6027.3.2.1.1.1.1.1.6.2 = STRING: "Gi 5/85 " << Channel member for Po2
dot3aCommonAggFdbIndex
SNMPv2-SMI::enterprises.6027.3.2.1.1.6.1.1.1107755009.1 = INTEGER: 1107755009
dot3aCommonAggFdbVlanId
SNMPv2-SMI::enterprises.6027.3.2.1.1.6.1.2.1107755009.1 = INTEGER: 1
dot3aCommonAggFdbTagConfig
SNMPv2-SMI::enterprises.6027.3.2.1.1.6.1.3.1107755009.1 = INTEGER: 2 (Tagged 1 or Untagged 2)
dot3aCommonAggFdbStatus
SNMPv2-SMI::enterprises.6027.3.2.1.1.6.1.4.1107755009.1 = INTEGER: 1 << Status active, 2 - status inactive
```

If we learn MAC addresses for the LAG, status is shown for those as well.

```
dot3aCurAggVlanId
SNMPv2-SMI::enterprises.6027.3.2.1.1.4.1.1.1.0.0.0.0.0.1.1 = INTEGER: 1
dot3aCurAggMacAddr
SNMPv2-SMI::enterprises.6027.3.2.1.1.4.1.2.1.0.0.0.0.0.1.1 = Hex-STRING: 00 00 00 00 00 00 01
dot3aCurAggIndex
SNMPv2-SMI::enterprises.6027.3.2.1.1.4.1.3.1.0.0.0.0.0.1.1 = INTEGER: 1
dot3aCurAggStatus
SNMPv2-SMI::enterprises.6027.3.2.1.1.4.1.4.1.0.0.0.0.0.1.1 = INTEGER: 1 <<
```

Status active, 2 - status
inactive

Layer 3 LAG does not include this support. SNMP trap works for the Layer 2 / Layer 3 / default mode LAG.

```
SNMPv2-MIB::sysUpTime.0 = Timeticks: (8500842) 23:36:48.42
SNMPv2-MIB::snmpTrapOID.0 = OID: IF-MIB::linkDown
IF-MIB::ifIndex.33865785 = INTEGER: 33865785
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "OSTATE_DN: Changed
interface state to down: Gi 0/0"
2010-02-10 14:22:39 10.16.130.4 [10.16.130.4]:
SNMPv2-MIB::sysUpTime.0 = Timeticks: (8500842) 23:36:48.42
SNMPv2-MIB::snmpTrapOID.0 = OID: IF-MIB::linkDown
IF-MIB::ifIndex.1107755009 = INTEGER: 1107755009
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "OSTATE_DN: Changed
interface state to down: Po 1"
2010-02-10 14:22:40 10.16.130.4 [10.16.130.4]:
SNMPv2-MIB::sysUpTime.0 = Timeticks: (8500932) 23:36:49.32 SNMPv2-
MIB::snmpTrapOID.0 = OID:
IF-MIB::linkUp IF-MIB::ifIndex.33865785 = INTEGER: 33865785 SNMPv2-
SMI::enterprises.6027.3.1.1.4.1.2 =
STRING: "OSTATE_UP: Changed interface state to up: Gi 0/0"
2010-02-10 14:22:40 10.16.130.4 [10.16.130.4]:
SNMPv2-MIB::sysUpTime.0 = Timeticks: (8500934) 23:36:49.34 SNMPv2-
MIB::snmpTrapOID.0 = OID:
IF-MIB::linkUp IF-MIB::ifIndex.1107755009 = INTEGER: 1107755009
SNMPv2-SMI::enterprises.6027.3.1.1.4.1.2 = STRING: "OSTATE_UP: Changed
interface state to up: Po 1"
```

BMP Functionality Using SNMP SET

To enable or disable BMP, use SNMP SET.

Table 66. List of BMP MIBS that have both read/write access

MIB Object	OID	Description
f10bmpMib	.1.3.6.1.4.1.6027.3.23	NODE
bmpEntry	.1.3.6.1.4.1.6027.3.23.1	NODE
bmpReloadType	.1.3.6.1.4.1.6027.3.23.1.1	LEAF INTEGER
bmpAutoSave	.1.3.6.1.4.1.6027.3.23.1.2	LEAF INTEGER
bmpConfigDownload	.1.3.6.1.4.1.6027.3.23.1.3	LEAF INTEGER
bmpDhcpTimeout	.1.3.6.1.4.1.6027.3.23.1.4	LEAF INTEGER
bmpRetryCount	.1.3.6.1.4.1.6027.3.23.1.5	LEAF INTEGER
bmpUserDefinedString	.1.3.6.1.4.1.6027.3.23.1.6	LEAF OCTET STRING

Entity MIBS

The Entity MIB provides a mechanism for presenting hierarchies of physical entities using SNMP tables.

The Entity MIB contains the following groups, which describe the physical elements and logical elements of a managed system. The following tables are implemented for the MXL switch.

Physical Entity

A physical entity or physical component represents an identifiable physical resource within a managed system.

Zero or more logical entities may utilize a physical resource at any given time. Determining which physical components are represented by an agent in the `EntPhysicalTable` is an implementation-specific matter. Typically, physical resources (for example, communications ports, backplanes, sensors, daughter-cards, power supplies, and the overall chassis), which can be managed via functions associated with one or more logical entities, are included in the MIB.

Containment Tree

Each physical component may be modeled as contained within another physical component.

A containment-tree is the conceptual sequence of `entPhysicalIndex` values that uniquely specifies the exact physical location of a physical component within the managed system. It is generated by following and recording each `entPhysicalContainedIn` instance up the tree towards the root, until a value of zero indicating no further containment is found.

Example of the Entity MIBS Outputs

```
Dell#show inventory optional-module
Unit Slot Expected   Inserted Next Boot Power
-----
0      0   QSFP+      QSFP+      AUTO Good
0      1  10GBASE-T  10GBASE-T  AUTO Good
1      0   QSFP+      QSFP+      AUTO Good
1      1  10GBASE-T  10GBASE-T  AUTO Good
2      0   QSFP+      QSFP+      AUTO Good
2      1   SFP+      SFP+      AUTO Good
```

The status for the MIBS is as follows:

```
vijayakrishnan@tapti[3:42pm] : /tftpboot > snmpwalk -c public -v 2c
10.16.130.135
1.3.6.1.2.1.47.1.1.1.1.2
SNMPv2-SMI::mib-2.47.1.1.1.1.2.1 = ""
SNMPv2-SMI::mib-2.47.1.1.1.1.2.2 = STRING: "PowerConnect MXL 10/40GbE"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.3 = STRING: "Module 0"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.4 = STRING: "Unit: 0 Port 1 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.5 = STRING: "Unit: 0 Port 2 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.6 = STRING: "Unit: 0 Port 3 10G Level"
SNMPv2-SMI::mib-2.47.1.1.1.1.2.7 = STRING: "Unit: 0 Port 4 10G Level"
```

```
SNMPv2-SMI::mib-2.47.1.1.1.1.2.8 = STRING: "Unit: 0 Port 5 10G Level"  
SNMPv2-SMI::mib-2.47.1.1.1.1.2.9 = STRING: "Unit: 0 Port 6 10G Level"  
---output truncated
```

Troubleshooting SNMP Operation

When you use SNMP to retrieve management data from an SNMP agent on a Dell Networking router, take into account the following behavior.

- When you query an IPv4 `icmpMsgStatsInPkts` object in the ICMP table by using the `snmpwalk` command, the output for echo replies may display incorrectly. To correctly display this information under ICMP statistics, use the `show ip traffic` command.
- When you query an `icmpStatsInErrors` object in the `icmpStats` table by using the `snmpget` or `snmpwalk` command, the output for IPv4 addresses may display incorrectly. To correctly display this information under IP and ICMP statistics, use the `show ip traffic` command.
- When you query an IPv4 `icmpMsgStatsInPkts` object in the ICMP table by using the `snmpwalk` command, the echo response output may not display. To correctly display ICMP statistics, such as echo response, use the `show ip traffic` command.

Stacking

Stacking is supported on the MXL switch platform.

Stacking is supported on a MXL 10/40GbE switch on the 40GbE ports (for the base module) or a 2-Port 40GbE QSFP+ module. You can connect up to six MXL 10/40GbE switches in a single stack. Stacking provides a single point of management and network interface controller (NIC) teaming for high availability and higher throughput.

Stacking MXL 10/40GbE Switches

A stack of MXL 10/40GbE switches operates as a virtual chassis with management units (primary and standby) and member units.

The Dell Networking operating software (OS) elects a primary (master) and secondary (standby) management unit; all other units are member units. The forwarding database resides on the master switch; all other stack units maintain a synchronized local copy. Each unit in the stack makes forwarding decisions based on their local copy.

The following example shows how you can stack four MXL 10/40GbE switches and the role played by each switch in the stack. The MXL 10/40GbE switches are connected to operate as a single stack in a ring topology using only the 40GbE ports on the base modules. You can use the 40GbE ports on the base module and FlexIO modules to create a stack in either a ring or daisy-chain topology.




Figure 117. Four-Stacked MXL 10/40GbE Switches

Stack Management Roles

The stack elects the management units for the stack management.

- **Stack master** — primary management unit, also called the master unit.
- **Standby** — secondary management unit.

The master holds the control plane and the other units maintain a local copy of the forwarding databases. From the stack master you can configure:

- System-level features that apply to all stack members.
- Interface-level features for each stack member.

The master synchronizes the following information with the standby unit:


- Stack unit topology
- Stack running configuration (which includes ACL, LACP, STP, SPAN, and so on.)
- Logs

The master switch maintains stack operation with minimal impact in the event of:

- Switch failure
- Inter-switch stacking link failure
- Switch insertion

- Switch removal

If the master switch goes off line, the standby replaces it as the new master and the switch with the next highest priority or MAC address becomes standby.

 **NOTE:** For the MXL switch, the entire stack has only one management IP address.

Stack Master Election

The stack elects a master and standby unit at bootup time based on two criteria.

- **Unit priority** – User-configurable. The range is from 1 to 14. A higher value (14) means a higher priority. The default is 0. To remove the stack-unit priority and set the priority back to the default value of zero, use the `no stack-unit priority` command.
- **MAC address (in case of priority tie)** – The unit with the higher MAC value becomes the master unit.

To view which switch is the stack master, enter the `show system` command. The following example shows sample output from an established stack.

A change in the stack master occurs when:

- You power down the stack master or bring the master switch offline.
- A failover of the master switch occurs.
- You disconnect the master switch from the stack.

When a stack reloads and all the units come up at the same time; for example, when all units boot up from flash, all units participate in the election and the master and standby are chosen based on the priority or MAC address.

When the units do not boot up at the same time, such as when some units are powered down just after reloading and powered up later to join the stack, they do not participate in the election process even though the units that boot up late may have a higher priority configured. This happens because the master and standby have already been elected, hence the unit that boots up late joins only as a member.

When an up and running standalone unit or stack is merged with another stack, based on election, the losing stack reloads and the master unit of the winning stack becomes the master of the merged stack.

For more details, see sections [Adding a Stack Unit](#) and [Merging Two Stacks](#). To ensure a fully synchronised bootup, it is possible to reset individual units to force them to give up the management role; or reload the whole stack from the command line interface (CLI).

Example of Viewing Stack Members

```
Dell# show system brief
```

```
Stack MAC : 00:1e:c9:f1:00:7b
```

```
Reload Type : jump-start [Next boot : normal-reload]
```

```
-- Stack Info --
```

Unit	UnitType	Status	ReqTyp	CurTyp	Version	Ports
0	Management	online	MXL-10/40GbE	MXL-10/40GbE	9-1-0-853	56
1	Standby	online	MXL-10/40GbE	MXL-10/40GbE	9-1-0-853	56
2	Member	online	MXL-10/40GbE	MXL-10/40GbE	9-1-0-853	56
3	Member	online	MXL-10/40GbE	MXL-10/40GbE	9-1-0-853	56


```
4 Member online MXL-10/40GbE MXL-10/40GbE 9-1-0-853 56
5 Member online MXL-10/40GbE MXL-10/40GbE 9-1-0-853 56
Dell#
```

Failover Roles

If the stack master fails (for example, is powered off), it is removed from the stack topology.

The standby unit detects the loss of peering communication and takes ownership of the stack management, switching from the standby role to the master role. The lack of a standby unit triggers an election within the remaining units for a standby role.

After the former master switch recovers, despite having a higher priority or MAC address, it does not recover its master role but instead takes the next available role.

MAC Addressing

All port interfaces in the stack use the MAC address of the management interface on the master switch.

The MAC address of the chassis in which the master MXL switch is installed is used as the stack MAC address.

The stack continues to use the master's chassis MAC address even after a failover. The MAC address is not refreshed until the stack is reloaded and a different unit becomes the stack master.

Stacking LAG

When multiple links are used between stack units, the Dell Networking OS automatically bundles them in a stacking LAG to provide aggregated throughput and redundancy.

The stacking LAG is established automatically and transparently by the Dell Networking OS (without user configuration) after peering is detected and behaves as follows:


- The stacking LAG dynamically aggregates; it can lose link members or gain new links.
- Shortest path selection inside the stack: If multiple paths exist between two units in the stack, the shortest path is used.

Supported Stacking Topologies

Stacking is supported on the MXL switch in ring and daisy-chain topologies.

Example 1: Dual-Ring Stack Across Multiple Chassis

Using two separate stacks in a dual-ring stacking topology provides redundancy and increased high availability in case of stack failure. Also, stacking upgrades are simplified when you have to take one stack offline, as shown in the following example.

 **NOTE:** A ring topology is recommended under normal operation because it provides increased resiliency when compared with a daisy chain topology. In daisy chain topology, any change in a non-edge stack unit causes a split stack.




Figure 118. Dual-Ring Stacking Topology for MXL 10/40GbE Switches

Example 2: Dual Daisy-Chain Stack Across Multiple Chassis

Using two separate, daisy-chained stacks in a stacking topology provides redundancy and increased high availability in case of stack failure. Also, stacking upgrades are simplified when you have to take one stack offline, as shown in the following example.




Figure 119. Dual Daisy-Chain Stacking Topology for MXL 10/40GbE Switches

Stack Group/Port Numbers

By default, each unit in Standalone mode is numbered stack-unit 0.

Stack-unit numbers are assigned to member switches when the stack comes up. The following example shows the stack-group numbers of 40GbE ports on an MXL 10/40GbE switch.




Figure 120. Stack-Group on an MXL 10/40GbE Switch

Configuring a Switch Stack

Configuring a switch stack is a four step process.

To configure and bring up a switch stack, follow these steps:

1. Connect the switches to be stacked with 40G direct attach or QSFP fibre cables.
2. Configure the stacking ports on each switch.
3. All switches must be booted together.
4. (Optional) Configure management priorities, unit numbers, or logical provisioning for stack units.

Stacking Prerequisites

Before you cable and configure a stack of MXL 10/40GbE switches, review the following prerequisites.

- All MXL 10/40GbE switches in the stack must be powered up with the initial or startup configuration before you attach the cables.
- All stacked MXL 10/40GbE switches must run the same Dell Networking OS version. The minimum Dell networking OS version required is 8.3.16.0. To check the version that a switch is running, use the `show version` command. To download a Dell Networking OS version, go to <http://www.dell.com/support/my-support/>.


- Stacking is supported only with other MXL 10/40GbE switches. A maximum of six MXL 10/40GbE switches is supported in a single stack. You cannot stack the MXL 10/40GbE switch with the M IO Aggregator or another type of switch.
- A maximum of four stack groups (40GbE ports) is supported on a stacked MXL 10/40GbE switch.
- Interconnect the stack units by following the instructions in [Cabling Stacked Switches](#).
- When you create stack ports on an MXL Switch, all ports must be fixed or on the expansion module. Mixing fixed and expansion module ports in order to stack is not supported.

Master Selection Criteria

A Master is elected or re-elected based on the following considerations, in order:

1. The switch with the highest priority at boot time.
2. The switch with the highest MAC address at boot time.
3. A unit is selected as Standby by the administrator, and a fail over action is manually initiated or occurs due to a Master unit failure.

No record of previous stack mastership is kept when a stack loses power. As it reboots, the election process will once again determine the Master and Standby switches. As long as the priority has not changed on any members, the stack will retain the same Master and Standby.

 **NOTE:** Each stack members' role (including the Master and Standby) can be defined by the user at any time by setting the priority.

If the entire stack is powered OFF and ON again, the unit that was the Master before the reboot will remain the Master after the stack resumes operation. However, when a stack is powered on, all members are in **sleep** mode for 5 seconds while waiting on the previous Master to join the stack. If the previous Master fails to join within 5 seconds, the remaining members (including the Standby) elect a new Master.

Configuring Priority and stack-group

Perform the following steps to configure the priorities and stack-groups for each of the switches.


1. Set the priorities for the stack-unit.

CONFIGURATION mode

```
stack-unit unit-number priority 1-14
```

```
Dell(conf)# stack-unit 0 priority 12
```

Setting the priority will determine which switch will become the management (Master) switch. The switch with the highest priority number is elected Master. The default priority is 0.

 **NOTE:** It is best practice to assign priority values to all switches before stacking them in order to acquire and retain complete control over each units role in the stack.

2. Configure the stack-group for each stack-unit.

CONFIGURATION mode


```
stack-unit unit-no stack-group stack-group-id
```

```
Dell(conf)# stack-unit 0 stack-group 0
```

```
Dell (conf)#02:26:05: %STKUNIT0-M:CP %IFMGR-6-STACK_PORTS_ADDED:
```

Ports Fo 0/33 have been configured as stacking ports.
Please save and reload for config to take effect

Stack-groups are easier to think of simply as stack ports. For example, using the `stack-group 0` command simply turns the lower port (port 9) into a stacking port. Similarly, `stack-group 1`, `stack-group 2` and `stack-group 3` commands correspond to ports 10, 11 and 12 respectively.

 **NOTE:** Stack-group is supported only in PMUX mode.

3. Continue to run the `stack-unit 0 stack-group <0-3>` command to add additional stack ports to the switch, using the stack-group mapping.

Cabling Stacked Switches

Before you configure MXL switches in a stack, connect the 40G direct attach or QSFP cables and transceivers to connect 40GbE ports on switches in the same or different chassis.

Cabling Restrictions

The following restrictions apply when setting up a stack of MXL 10/40GbE switches.

- Only daisy-chain or ring topologies are supported; star and full mesh topologies are not supported.
- Stacking is supported only on 40GbE links by connecting 40GbE ports on the base module or a 2-Port QSFP+ module. Stacking is not supported on 10GbE ports or 4x10GbE ports.
 - To convert the 40GbE ports on the 2-Port QSFP+ module from their default 4x10GbE mode of operation to 40GbE mode, refer to [Converting 4x10GbE Ports to 40GbE for Stacking](#).
- Use only QSFP transceivers and QSFP or direct attach cables (purchased separately) to connect stacking ports.

Cabling Redundancy

Connect the units in a stack with two or more stacking cables to avoid a stacking port or cable failure. Removing one of the stacked cables between two stacked units does not trigger a reset.

Cabling Procedure

The following cabling procedure uses the stacking topology shown earlier in this chapter.

Follow the same steps to cable switches in any of the stacking topologies shown in [Supported Stacking Topologies](#).

1. Connect a 40GbE port on the first switch to a 40GbE port on the second switch.
2. Connect another 40GbE port on the second switch to a 40GbE port on the third switch.
3. Connect another 40GbE port on the third switch to a 40GbE port on the fourth switch.
4. Connect another 40GbE port on the fourth switch to a 40GbE port on the first switch.

The resulting ring topology allows the entire stack to function as a single switch with resilient fail-over capabilities. If you do not connect the last switch to the first switch (Step 4), the stack operates in a daisy chain topology with less resiliency. Any failure in a non-edge stack unit causes a split stack.


Accessing the CLI

To configure a stack, you must access the stack master in one of the following ways.

- For remote out-of-band management (OOB), enter the OOB management interface IP address into a Telnet or secure shell (SSH) client and log in to the switch using the user ID and password to access the CLI.
- For local management, use the attached console connection to the master switch to log in to the CLI. Console access to the stack CLI is available on the master only.
- For remote in-band management from a network management station, enter the virtual local area network (VLAN) IP address of the management port and log in to the switch to access the CLI.

Configuring and Bringing Up a Stack

After you attach the 40G QSFP or direct attach cables in a stack of MXL 10/40GbE Switches, to bring up the stack, follow these steps.

 **NOTE:** The procedure uses command examples for the stacking topology shown previously in this chapter.

1. Set up a connection to the CLI on an MXL 10/40GbE Switch as described in [Accessing the CLI](#).
2. Log on to the CLI and enter Global Configuration mode.

Login: username

Password: *****

Dell> enable

Dell# configure

3. Configure a 40GbE port for stacking mode.

CONFIGURATION mode

```
stack-unit unit-number stack-group group-number
```

The valid values are from 0 to 5.

The default value is **0**.

- `stack-unit <unit-number>`: is the unit-number of the member stack unit.
- `stack-group group-number` is the number of stacked port on unit. The valid values are from 0 to 1.

4. Save the stacking configuration on the 40GbE ports.


EXEC PRIVILEGE mode

```
write memory
```

5. Repeat Steps 1 to 4 on each MXL 10/40GbE switch in the stack.
6. Log on to the CLI and reboot each switch, one after another, in as short a time as possible.

EXEC PRIVILEGE mode

```
reload
```

 **NOTE:** If the stacked switches all reboot at approximately the same time, the switch with the highest MAC address is automatically elected as the master switch. The switch with the next highest MAC address is elected as standby. As each switch joins the stack, it is assigned the lowest available stack-unit number from 0 to 5. The default configuration of each stacked switch is stored in the running configuration of the stack. The stack-unit ID numbers are retained after future stack reloads.

To verify the stack-unit number assigned to each switch in the stack, use the `show system brief` command.

To configure stacked switches so that stacking roles are determined by preset priorities, use the `stack-unit priority` command (refer to [Assigning a Priority to Stacked Switches](#)).

Assigning a Priority to Stacked Switches

To configure the stack so that the roles are assigned according to pre-determined priorities instead of using the highest MAC addresses, use the `stack-unit priority` command in Global Configuration mode on each stacked switch.

The switch with the highest priority number is elected master. The switch with the next highest priority number is elected standby and takes over stack management if the master switch fails.

Configure the priority of stacked switches to determine stack mastership.
CONFIGURATION mode

```
stack-unit unit-number priority number
```

- `stack-unit unit-number` identifies the switch in the stack.
- `priority priority-number` specifies the management priority.

The valid range is from 1 to 14.

The default is **0**.

To revert the management priority of a stack unit to the default value of 0, use the `no` form of the `stack-unit unit-number priority number` command.

After you reconfigure the priorities of stacked switches, reload the stack so that a new master and standby election is performed.

Renumbering a Stack Unit

To renumber a stack unit to reset the unit numbering for a master, standby or member unit, use the following command.

- If you renumber the master switch, you are prompted to reload the entire stack.
- If you renumber the standby switch, only the switch reloads and is replaced by a member switch that is elected as the new standby.
- If you renumber a member switch, only the member switch reloads.

- If you renumber a switch to a number already assigned to another stack unit, the following error message displays:

```
Dell#stack-unit 5 renumber 0
% ERROR: stack unit 0 already exists.
```

Assign a stack-number to a unit.

EXEC Privilege mode

```
stack-unit unit-number renumber new-number
```

Provisioning a Stack Unit

You can logically provision a stack-unit number to accept only an MXL 10/40GbE switch.

Provisioning is a type of pre-configuration that is stored on the master switch and applied when a stacked unit is assigned the unit number.

When you provision a unit number for an MXL 10/40GbE switch:

- The base-module ports on the switch (ports 33 and 37/stack groups 0 and 1) are pre-configured for 40GbE operation.
- The 40GbE ports on FlexIO modules (ports 41 and 45 in slot 0; ports 49 and 53 in slot 1) are pre-configured for 4x10GbE (quad mode) operation.

Create a virtual stack unit by logically provisioning a switch.

CONFIGURATION mode

```
stack-unit unit-number provision MXL-10/40GbE
```

To provision a stack unit, use the `stack-unit provision` command in Global Configuration mode, save the provisioning configuration, and reload the stack.

Dell Networking OS Behavior: Stacking configuration is handled as follows on an MXL 10/40GbE switch:

- If a stack unit goes down and is removed from the stack, the logical provisioning configured for the stack-unit number is saved on the master and standby switches.
- When you add a new unit to the stack and the stack already has an existing member unit with the same stack-unit number, the new unit is assigned the smallest available unit number (from 0 to 5). A configuration mismatch between the newly added unit and a logically provisioned unit occurs in the following situations:
 - The logical provisioning for the unit number configures FlexIO module ports for 4x10GbE operation and the added unit has FlexIO Module ports operating in 40GbE mode.
 - The logical provisioning for the unit number and the added unit have different stack groups configured.
- When a configuration mismatch occurs, the newly added switch enters into a Card-Problem state and is disabled. A syslog error message generates. To restore a stacked switch in a Card-Problem state, refer to [Failure Scenarios](#).
- A stack unit can also enter a Card-Problem state after a split-stack reload in which a unit that was previously neither the master nor standby is elected as the new master and has logical stack-unit provisioning configured for a stack-unit number that creates a mismatch with the stack-unit numbering on other units.

Converting 4x10GbE Ports to 40GbE for Stacking

Stacking is supported only on 40GbE links by connecting 40GbE ports on the base module or a 2-Port QSFP+ module.

However, on a 2-Port 40GbE QSFP+ module, the ports operate by default in 4x10GbE (quad) mode with breakout cables as eight 10GbE ports.

Change a port from 4x10GbE to 40GbE mode of operation for stacking.

```
no stack-unit port portmode quad
```

After you convert the 4x10GbE ports to 40GbE, you must save the configuration and reload the stack for the change to take effect.

- `stack-unit unit-number` is the unit ID number in the stack unit. The valid range is from 0 to 5.
- `port port-number` specifies the port number of the QSFP+ port to be converted to 40GbE mode. The valid values are: base-module ports: 33 or 37; slot 0: 41 or 45; slot 1: 49 or 53. As shown below, `portmode quad` identifies the port as a split 10GbE SFP+ port.

Example of the `no stack-unit port portmode quad` Command

```
Dell(conf)# no stack-unit unit-number port port-number portmode quad
```

To display the stack-unit number, use the `show system brief` command.

Removing a Port from the Stacking Mode

To remove a 40GbE port from the stack, use the `no` form of the `stack-unit unit-number stack-group number` command.

After entering the command, save the configuration and reload the stack for the change to take effect.

Remove a stacked port from a stack.

CONFIGURATION mode

```
no stack-unit unit-number stack-group group end write memory reload
```

When the reload completes, the port comes up in 40GbE mode if it is on the base module and in 4x10GbE (quad) mode if the port is on a FlexIO module, such as a 2-Port 40GbE QSFP+ module.

Removing a Switch from a Stack

To remove a switch from a stack, disconnect the stacking cables from the unit either when the unit is powered on or off and is online or offline.

After you remove all 40GbE ports from a stack ([Removing a Port from the Stacking Mode](#)), the switch functions in standalone mode but retains the running and startup configuration that was last synchronized by the master switch while it operated as a stack unit.

If you remove a unit from the middle of a stack, the stack is split into multiple parts. Each split stack forms a new stack according to MAC addresses or assigned priorities, as described in [Configuring and Bringing Up a Stack](#) and [Assigning a Priority to Stacked Switches](#).

Adding a Stack Unit

You can add a new unit to an existing stack both when the unit has no stacking ports (stack groups) configured and when the unit already has stacking ports configured.

If the units to be added to the stack have been previously used, they are assigned the smallest available unit ID in the stack.

If a standalone switch has no stack groups configured, you can add it to a stack. To add a standalone switch to a stack, follow these steps.

1. Power on the switch.
2. Attach QSFP or direct attach cables to connect 40GbE ports on the switch to one or more switches in the stack.
3. Log on to the CLI and enter global configuration mode.

```
Login: username
```

```
Password: *****
```

```
Dell> enable
```

```
Dell# configure
```

4. Configure a 40GbE port for stacking.

```
CONFIGURATION mode
```

```
stack-unit 0 stack-group group-number
```

- `stack-unit 0` defines the default ID unit-number in the initial configuration of a switch.
- `stack-group group-number` configures a 40GbE port for stacking. Base-module ports are stack groups 0 and 1; 40GbE ports on a FlexIO module in slot 0 are stack groups 2 and 3 and in slot 1 are stack groups 4 and 5.

5. Save the stacking configuration on the 40GbE ports.

```
EXEC Privilege mode
```

```
write memory
```

6. Reload the switch.

```
EXEC Privilege mode
```

```
reload
```

The Dell Networking OS automatically assigns a number to the new unit and adds it as member switch in the stack. The new unit synchronizes its running and startup configurations with the stack.

If a standalone switch already has stack groups configured, continue with these steps:

7. Attach QSFP or direct attach cables to connect the 40GbE ports already configured as stack groups on the switch to one or more switches in the stack.
8. Power on the switch. The Dell Networking OS automatically assigns a number to the new unit and adds it as member switch in the stack. The new unit synchronizes its running and startup configurations with the stack.

Dell Networking OS Behavior: When you add a new switch to a stack:


- If the new unit has been configured with a stack number that is already assigned to a stack member, the stack avoids a numbering conflict by assigning the new switch the first available stack number.
- If the stack has been provisioned for the stack number that is assigned to the new unit, the pre-configured provisioning must match the switch type. If there is a conflict between the provisioned switch type and the new unit, a mismatch error message displays.

Merging Two Stacks

You can merge two MXL 10/40GbE Switch stacks while they are powered and online.

To merge two stacks, connect one stack to the other with 40G QSFP or direct attach cables. After you connect the stacking cables, a merge of the two stacks is performed:

- The Dell Networking OS selects a master switch for the merged stack from the existing masters in the two stacks. To ensure that one of the two master switches wins the master election in the merged stack, use the `stack-unit priority` command to configure the highest priority for the unit (refer to [Assigning a Priority to Stacked Switches](#)).
- All the units in the losing stack go for a reboot and then merge with the winning stack that has the stack master.
- If there is no unit numbering conflict, the stack members retain their previous unit numbers. Otherwise, the stack master assigns new unit numbers, based on the order in which they come online.
- The new stack master uses its own startup and running configurations to synchronize the configurations on the new stack members.

 **NOTE:** Adding a new unit that is powered on and has stack groups configured is the same as merging two stacks (refer to [Adding a Stack Unit](#)). If the new unit has been configured with a higher priority than the current stack master, it becomes the new stack master and the stack reloads. If the new unit does not have a higher priority than the master switch, it is added as a member switch.

Splitting a Stack

To split an MXL 10/40GbE switch stack, unplug the stacking cables between member units at any time: while the stack is powered on or off and when the units are online or offline.

Each portion of the split stack retains the startup and running configuration of the original stack.

For a stack that is split into two smaller stacks, each with multiple units:

- If one of the new stacks receives the master and the standby units, it is unaffected by the split.
- If one of the new stacks receives only the master unit, the master switch retains its role and a new standby is elected.
- If one of the new stacks receives only the standby unit, it becomes the master in the new stack and the Dell Networking OS elects a new standby.
- If one of the new stacks receives neither the master nor the standby unit, the stack is reset so that a new election takes place.

Managing Redundant Stack Management

To manage the redundancy behavior in a stack, use the following `redundancy` commands.

- Reset the current stack master and make the standby unit the new master.
EXEC Privilege mode

```
redundancy force-failover stack-unit
```

A new standby is elected. When the former stack master comes back online, it becomes a member unit.

- Prevent the stack master from rebooting after a failover.
CONFIGURATION mode

```
redundancy disable-auto-reboot stack-unit
```

This command does not affect a forced failover, manual reset, or a stack-link disconnect.

- Display redundancy information.
EXEC Privilege mode

```
show redundancy
```

Resetting a Unit on a Stack

To reload any of the member units or the standby in a stack, use the following reset commands. If you try to reset the stack master, an error message displays:

```
% Error: Reset of master unit is not allowed.
```

To rest a unit on a stack, use the following commands.

- Reload a stack-unit.
EXEC Privilege mode

```
reset stack-unit unit-number
```

- Reload a member unit, from the unit itself.
EXEC Privilege mode

```
reset-self
```

- Reset a stack-unit when the unit is in a problem state.
EXEC Privilege mode

```
reset stack-unit unit-number hard
```

Verify a Stack Configuration

The following lists the status of a stacked switch according to the color of the System Status light emitting diodes (LEDs) on its front panel.

- Blue indicates the switch is operating as the stack master or as a standalone unit.
- Off indicates the switch is a member or standby unit.
- Amber indicates the switch is booting or a failure condition has occurred.

Using Show Commands

To display information on the stack configuration, use the `show` commands on the master switch.

- Displays stacking roles (master, standby, and member units) and the stack MAC address.
`show system [brief]`
- Displays the FlexIO modules currently installed in expansion slots 0 and 1 on a switch and the expected module logically provisioned for the slot.
`show inventory optional-module`
- Displays the stack groups allocated on a stacked switch. The valid stack-unit numbers are from 0 to 5.
`show system stack-unit unit-number stack-group configured`
- Displays the port numbers that correspond to the stack groups on a switch. The valid stack-unit numbers are from 0 to 5.
`show system stack-unit unit-number stack-group`
- Displays the type of stack topology (ring or daisy chain) with a list of all stacked ports, port status, link speed, and peer stack-unit connection.
`show system stack-ports [status | topology]`

Example of the `show system brief` Command

Example of the `show system` Command

Example of the `show inventory optional-module` Command

Example of the `show system stack-unit stack-group configured` Command

Example of the `show system stack-unit stack-group` Command

Example of the `show system stack-ports (ring)` Command

Example of the `show system stack-ports (daisy chain)` Command

```
Dell# show system brief
```

```
Stack MAC : 00:1e:c9:f1:00:7b
```

```
Reload Type : jump-start [Next boot : normal-reload]
```

```
-- Stack Info --
```

Unit	UnitType	Status	ReqTyp	CurTyp	Version	Ports
0	Management	online	MXL-10/40GbE	MXL-10/40GbE	9-1-0-853	56
1	Standby	online	MXL-10/40GbE	MXL-10/40GbE	9-1-0-853	56
2	Member	online	MXL-10/40GbE	MXL-10/40GbE	9-1-0-853	56
3	Member	online	MXL-10/40GbE	MXL-10/40GbE	9-1-0-853	56
4	Member	online	MXL-10/40GbE	MXL-10/40GbE	9-1-0-853	56
5	Member	online	MXL-10/40GbE	MXL-10/40GbE	9-1-0-853	56

```
Dell#show system
```

```
Stack MAC : 00:1e:c9:f1:00:e3
```

```
Reload Type : normal-reload [Next boot : normal-reload]
```

```
-- Unit 0 --
```

```
Unit Type      : Member Unit  
Status         : not present  
Required Type  : MXL-10/40GbE - 34-port GE/TE/FG (XL)
```

```

-- Unit 1 --
Unit Type      : Management Unit
Status        : online
Next Boot     : online
Required Type  : MXL-10/40GbE - 34-port GE/TE/FG (XL)
Current Type   : MXL-10/40GbE - 34-port GE/TE/FG (XL)
Master priority : 14
Hardware Rev   : 2.0
Num Ports     : 56
Up Time       : 19 hr, 30 min
Dell Networking OS Version : 9-1-0-1010
Jumbo Capable : yes
POE Capable    : no
Burned In MAC  : 00:1e:c9:f1:00:e3
No Of MACs    : 3

```

```

-- Unit 2 --
Unit Type      : Member Unit
Status        : online
Next Boot     : online
Required Type  : MXL-10/40GbE - 34-port GE/TE/FG (XL)
Current Type   : MXL-10/40GbE - 34-port GE/TE/FG (XL)
Master priority : 12
Hardware Rev   : 2.0
Num Ports     : 56
Up Time       : 19 hr, 30 min
Dell Networking OS Version : 9-1-0-1010
Jumbo Capable : yes
POE Capable    : no
Burned In MAC  : 00:1e:c9:f1:00:c7
No Of MACs    : 3

```

```

-- Unit 3 --
Unit Type      : Member Unit
Status        : not present
Required Type  : MXL-10/40GbE - 34-port GE/TE/FG (XL)

```

```

-- Unit 4 --
Unit Type      : Standby Unit
Status        : online
Next Boot     : online
Required Type  : MXL-10/40GbE - 34-port GE/TE/FG (XL)
Current Type   : MXL-10/40GbE - 34-port GE/TE/FG (XL)
Master priority : 13
Hardware Rev   : 3.0
Num Ports     : 56
Up Time       : 19 hr, 30 min
Dell Networking OS Version : 9-1-0-1010
Jumbo Capable : yes
POE Capable    : no

```

```
Dell# show inventory optional-module
```

```

Unit Slot Expected Inserted Next Boot Power
-----
0    0    SFP+    SFP+    AUTO    Good
0    1    QSFP+    QSFP+    AUTO    Good
* - Mismatch

```

```
Dell# show system stack-unit 1 stack-group configured
Configured stack groups in stack-unit 1
```

```

-----
0
1
4
5

Dell#show system stack-unit 1 stack-group
Stack group      Ports
-----
      0          0/33
      1          0/37
      2          0/41
      3          0/45
      4          0/49
      5          0/53

Dell#

```

```

Dell# show system stack-ports
Topology: Ring
Interface Connection Link Speed Admin Link Trunk
              (Gb/s)  Status Status Group
0/33         1/37      40      up   up
0/37         2/33      40      up   up
0/41         1/49      40      up   up
0/45         2/53      40      up   up
1/33         2/37      40      up   up
1/37         0/33      40      up   up
1/49         0/41      40      up   up
1/53         2/49      40      up   up
2/33         0/37      40      up   up
2/37         1/33      40      up   up
2/49         1/53      40      up   up
2/53         0/45      40      up   up

```

```

Dell# show system stack-ports
Topology: Daisy Chain

Interface Connection Link Speed Admin Link Trunk
              (Gb/s)  Status Status Group
0/33         1/37      40      up   up
0/41         1/49      40      up   up
1/33         2/37      40      up   up
1/37         0/33      40      up   up
1/49         0/41      40      up   up
1/53         2/49      40      up   up
2/37         1/33      40      up   up
2/49         1/53      40      up   up

```

Troubleshooting a Switch Stack

To perform troubleshooting operations on a switch stack, use the following commands on the master switch.

1. Displays the status of stacked ports on stack units.
`show system stack-ports`
2. Displays the master standby unit status, failover configuration, and result of the last master-standby synchronization; allows you to verify the readiness for a stack failover.


```
show redundancy
```

3. Displays input and output flow statistics on a stacked port.

```
show hardware stack-unit unit-number stack-port port-number
```

4. Clears statistics on the specified stack unit. The valid stack-unit numbers are from 0 to 5.

```
clear hardware stack-unit unit-number counters
```

Example of the show system stack-ports Command

Example of the show redundancy Command

Example of the show hardware stack-unit port-stack Command

```
Dell# show system stack-ports
Topology: Ring
```

Interface	Connection	Link Speed (Gb/s)	Admin Status	Link Status	Trunk Group
0/41	2/45	40	up	up	
0/45	1/41	40	up	up	
1/41	0/45	40	up	up	
1/45	2/41	40	up	up	
2/41	1/45	40	up	up	
2/45	0/41	40	up	up	

```
Dell#show redundancy
```

```
-- Stack-unit Status --
```

```
-----
Mgmt ID:                0
Stack-unit ID:          0
Stack-unit Redundancy Role: Primary
Stack-unit State:       Active (Indicates Master Unit.)
Stack-unit SW Version:  E8-3-16-79
Link to Peer:           Up
```

```
-- PEER Stack-unit Status --
```

```
-----
Stack-unit State:       Standby (Indicates Standby Unit.)
Pe er stack-unit ID:   2
Stack-unit SW Version: E8-3-16-79
```

```
-- Stack-unit Redundancy Configuration --
```

```
-----
Primary Stack-unit:    mgmt-id 0
Auto Data Sync:        Full
Failover Type:         Hot (Failover Failover type with redundancy.)
force-failover
Auto reboot Stack-unit: Disabled
Auto failover limit:   3 times in 60 minutes
```

```
-- Stack-unit Failover Record --
```

```
-----
Failover Count:        0
Last failover timestamp: None
Last failover Reason:  None
Last failover type:    None
```

```
-- Last Data Block Sync Record: --
```

```
-----
Stack Unit Config:     succeeded Mar 24 2012 20:07:39
Start-up Config:       succeeded Mar 24 2012 20:07:39 (Latest sync of config.)
Runtime Event Log:     succeeded Mar 24 2012 20:07:39
```

```
Running Config:      succeeded Mar 24 2012 20:07:39
ACL Mgr:            succeeded Mar 24 2012 20:07:39
LACP: no block sync done
STP: no block sync done
```

```
Dell# show hardware stack-unit 1 stack-port 53
Input Statistics:
 7934 packets, 1049269 bytes
 0 64-byte pkts, 7793 over 64-byte pkts, 100 over 127-byte pkts
 0 over 255-byte pkts, 7 over 511-byte pkts, 34 over 1023-byte pkts
 70 Multicasts, 0 Broadcasts
 0 runts, 0 giants, 0 throttles
 0 CRC, 0 overrun, 0 discarded
Output Statistics:
 438 packets, 270449 bytes, 0 underruns
 0 64-byte pkts, 57 over 64-byte pkts, 181 over 127-byte pkts
 54 over 255-byte pkts, 0 over 511-byte pkts, 146 over 1023-byte pkts
 72 Multicasts, 0 Broadcasts, 221 Unicasts
 0 throttles, 0 discarded, 0 collisions, 0 wredDrops
Rate info (interval 45 seconds):
 Input 00.00 Mbits/sec,    0 packets/sec, 0.00% of line-rate
 Output 00.00 Mbits/sec,   0 packets/sec, 0.00% of line-rate
```

Failure Scenarios

The following sections describe some of the common fault conditions that can happen in a switch stack and how they are resolved.

Stack Member Fails

- **Problem:** A unit that is not the stack master fails in an operational stack.
- **Resolution:** If a stack member fails in a daisy chain topology, a split stack occurs. If a member unit fails in a ring topology, traffic is re-routed over existing stack links.

The following syslog messages are generated when a member unit fails:

```
Dell#May 31 01:46:17: %STKUNIT3-M:CP %IPC-2-STATUS: target stack unit 4 not responding
```

```
May 31 01:46:17: %STKUNIT3-M:CP %CHMGR-2-STACKUNIT_DOWN: Major alarm: Stack unit 4 down - IPC timeout
```

```
Dell#May 31 01:46:17: %STKUNIT3-M:CP %IFMGR-1-DEL_PORT: Removed port: Te 4/1-32,41-48, Fo 4/49,53
```

```
Dell#May 31 01:46:18: %STKUNIT5-S:CP %IFMGR-1-DEL_PORT: Removed port: Te 4/1-32,41-48, Fo 4/49,53
```

Unplugged Stacking Cable

- **Problem:** A stacking cable is unplugged from a member switch. The stack loses half of its bandwidth from the disconnected switch.
- **Resolution:** Intra-stack traffic is re-routed on a another link using the redundant stacking port on the switch. A recalculation of control plane and data plane connections is performed.

Master Switch Fails

- **Problem:** The master switch fails due to a hardware fault, software crash, or power loss.
- **Resolution:** A failover procedure begins:
 1. Keep-alive messages from the MXL 10/40GbE master switch time out after 60 seconds and the switch is removed from the stack.
 2. The standby switch takes the master role. Data traffic on the new master switch is uninterrupted. Protocol traffic is managed by the control plane.
 3. A member switch is elected as the new standby. Data traffic on the new standby is uninterrupted. The control plane prepares for operation in Warm Standby mode.

Stack-Link Flapping Error

Problem/Resolution: Stacked MXL 10/40GbE Switches monitor their own stack ports and disable any stack port that flaps five times within 10 seconds. If the stacking ports that flap are on the master or standby, KERN-2-INT error messages note the units

To re-enable a downed stacking port, power cycle the stacked switch on which the port is installed.

The following is an example of the stack-link flapping error message.

```
-----MANAGEMENT
UNIT-----
Error: Stack Port 49 has flapped 5 times within 10 seconds. Shutting down this
stack port now.
Error: Please check the stack cable/module and power-cycle the stack.
10:55:20: %STKUNIT1-M:CP %KERN-2-INT: Error: Stack Port 50 has flapped 5 times
within 10
seconds. Shutting down this stack port now.
10:55:20: %STKUNIT1-M:CP %KERN-2-INT: Error: Please check the stack cable/
module and
power-cycle the stack.
-----STANDBY
UNIT-----
10:55:18: %STKUNIT1-M:CP %KERN-2-INT: Error: Stack Port 50 has flapped 5 times
within 10
seconds. Shutting down this stack port now.
10:55:18: %STKUNIT1-M:CP %KERN-2-INT: Error: Please check the stack cable/module
and power-cycle the stack.
-----MEMBER
2-----
Error: Stack Port 51 has flapped 5 times within 10 seconds. Shutting down this
stack port now.
Error: Please check the stack cable/module and power-cycle the stack.
```

Master Switch Recovers from Failure

- **Problem:** The master switch recovers from a failure after a reboot and rejoins the stack:
 - As a member unit if there is already a standby
 - As a standby if there is no standby in the stack

The protocol and control plane recovery requires time before the switch is fully online.

- **Resolution:** When the entire stack is reloaded, the recovered master switch becomes the master unit of the stack.

Stack Unit in Card-Problem State Due to Incorrect Dell Networking OS Version

- **Problem:** A stack unit enters a Card-Problem state because the switch has a different the Dell Networking OS version than the master unit. The switch does not come online as a stack unit.
- **Resolution:** To restore a stack unit with an incorrect the Dell Networking OS version as a member unit, disconnect the stacking cables on the switch and install the correct the Dell Networking OS version. Then add the switch to the stack as described in [Adding a Stack Unit](#). To verify that the problem has been resolved and the stacked switch is back online, use the `show system brief` command.

```
Dell#show system brief
Stack MAC : 00:1e:c9:f1:01:57
Reload Type : normal-reload [Next boot : normal-reload]
```

```
-- Stack Info --
Unit UnitType  Status          ReqTyp          CurTyp Version Ports
-----
0  Management online          MXL-10/40GbE   MXL-10/40GbE   8-3-16-79 56
1  Member      card problem MXL-10/40GbE           unknown      56
2  Standby    online          MXL-10/40GbE   MXL-10/40GbE   8-3-16-79 56
3  Member    not present
4  Member    not present
5  Member    not present
```

Card Problem — Resolved

```
Dell#show system brief
Stack MAC : 00:1e:c9:f1:01:57
Reload Type : normal-reload [Next boot : normal-reload]
```

```
-- Stack Info --
Unit UnitType  Status          ReqTyp          CurTyp Version Ports
-----
0  Management online          MXL-10/40GbE   MXL-10/40GbE   8-3-16-79 56
1  Member      online        MXL-10/40GbE           unknown      56
2  Standby    online          MXL-10/40GbE   MXL-10/40GbE   8-3-16-79 56
3  Member    not present
4  Member    not present
5  Member    not present
```

Stack Unit in Card-Problem State Due to Configuration Mismatch

- **Problem:** A stack unit enters a Card-Problem state because there is a configuration mismatch between the logical provisioning stored for the stack-unit number on the master switch and the newly added unit with the same number.
- **Resolution:** The resolution is to reload the stack. When the stack is up, the card problem will be solved.

To correct a configuration mismatch, reload the entire stack using the `reload` command in EXEC Privilege mode.

Upgrading a Switch Stack

To upgrade all switches in a stack with the same Dell Networking OS version, follow these steps.

1. Copy the new Dell Networking OS image to a network server.
2. Download the Dell Networking OS image by accessing an interactive CLI that requests the server IP address and image filename, and prompts you to upgrade all member stack units.
EXEC Privilege mode

```
'upgrade system { flash: | ftp: | scp: | tftp: | usbflash: } partition
```

Specify the system partition on the master switch into which you want to copy the Dell Networking OS image. The system then prompts you to upgrade all member units with the new Dell Networking OS version.

The valid values are a: and b:.

3. Reboot all stack units to load the Dell Networking OS image from the same partition on all switches in the stack.
CONFIGURATION mode

```
boot system stack-unit all primary system partition
```

4. Save the configuration.
CONFIGURATION mode

```
write memory
```

5. Reload the stack unit to activate the new Dell Networking OS version.
CONFIGURATION mode

```
reload
```

Example of Upgrading all Stacked Switches

The following example shows how to upgrade all switches in a stack, including the master switch.

```
Dell# upgrade system ftp: A:
Address or name of remote host []: 10.11.200.241
Source file name []: $V-9-1-0/NAVASOTA-DEV-9-1-0-887/Dell-XL-9-1-0-887.bin
User name to login remote host: ftp
Password to login remote host:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Erasing IOM Primary Image, please wait
.!.....
...
.....Writing.....
...
...
...
...
31972272 bytes successfully copied
System image upgrade completed successfully.
```

```

Upgrade system image for all stack-units [yes/no]: yes
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!
Image upgraded to all
Dell# configure
Dell(conf)# boot system stack-unit all primary system: A:
Dell(conf)# end
Dell# write memory
Jan 3 14:01:48: %STKUNIT0-M:CP %FILEMGR-5-FILESAVED: Copied running-config to
startup-config
in flash by default
Synchronizing data to peer Stack-unit
!!!!
Dell# reload
Proceed with reload [confirm yes/no]: yes

```

Upgrading a Single Stack Unit

Upgrading a single stacked switch is necessary when the unit was disabled due to an incorrect Dell Networking OS version.

This procedure upgrades the image in the boot partition of the member unit from the corresponding partition in the master unit.

1. Download the Dell Networking OS image from the master's boot partition to the member unit, and upgrade the relevant boot partition in the single stack-member unit.
EXEC Privilege mode

```
upgrade system stack-unit unit-number partition
```

2. Reboot the stack unit from the master switch to load the Dell Networking OS image from the same partition.
CONFIGURATION mode

```
boot system stack-unit unit-number primary system partition
```

3. Save the configuration.
EXEC Privilege mode

```
write memory
```

4. Reset the stack unit to activate the new Dell Networking OS version.
EXEC Privilege mode

```
power-cycle stack-unit unit-number
```

Example of Upgrading a Single Stack Unit

The following example shows how to upgrade an individual stack unit.

```

Dell# upgrade system stack-unit 2 A:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Image upgraded to Stack unit 2

```

```
Dell# configure
Dell(conf)# boot system stack-unit 2 primary system: A:
Dell(conf)# end
Dell#Jan 3 14:27:00: %STKUNIT0-M:CP %SYS-5-CONFIG_I: Configured from console
Dell# write memory
Jan 3 14:27:10: %STKUNIT0-M:CP %FILEMGR-5-FILESAVED: Copied running-config to
startup-config in flash
by default
Synchronizing data to peer Stack-unit
!!!!
....
Dell# power-cycle stack-unit 2
Proceed with power-cycle? Confirm [yes/no]:yes
```

Storm Control

Storm control is supported on the MXL switch platform.

The storm control feature allows you to control unknown-unicast and broadcast traffic on Layer 2 and Layer 3 physical interfaces.

Dell Networking OS Behavior: The Dell Networking OS supports broadcast control (the `storm-control broadcast` command) for Layer 2 and Layer 3 traffic. The minimum number of packets per second (PPS) that storm control can limit is two.

Configure Storm Control

Storm control is supported in INTERFACE mode and CONFIGURATION mode.

Configuring Storm Control from INTERFACE Mode

To configure storm control, use the following command.

You can only configure storm control for ingress traffic in INTERFACE mode. If you configure storm control from both INTERFACE and CONFIGURATION mode, the INTERFACE mode configurations override the CONFIGURATION mode configurations.

- Configure storm control.
INTERFACE mode

```
storm control
```

Configuring Storm Control from CONFIGURATION Mode

To configure storm control from CONFIGURATION mode, use the following command.

You can configure storm control for ingress traffic in CONFIGURATION mode. Do not apply per-virtual local area network (per-VLAN) quality of service (QoS) on an interface that you have enabled storm-control (either on an interface or globally).

- Configure storm control.
CONFIGURATION mode

```
storm control
```


Spanning Tree Protocol (STP)

The spanning tree protocol (STP) is supported on the MXL switch platform.

Protocol Overview

STP is a Layer 2 protocol — specified by IEEE 802.1d — that eliminates loops in a bridged topology by enabling only a single path through the network.

By eliminating loops, the protocol improves scalability in a large network and allows you to implement redundant paths, which can be activated after the failure of active paths. Layer 2 loops, which can occur in a network due to poor network design and without enabling protocols like xSTP, can cause unnecessarily high switch CPU utilization and memory consumption.

The Dell Networking OS supports three other variations of spanning tree, as shown in the following table.

Table 67. Dell Networking OS Supported Spanning Tree Protocols

Dell Networking Term	IEEE Specification
Spanning Tree Protocol (STP)	802.1d
Rapid Spanning Tree Protocol (RSTP)	802.1w
Multiple Spanning Tree Protocol (MSTP)	802.1s
Per-VLAN Spanning Tree Plus (PVST+)	Third Party

Configure Spanning Tree

Configuring spanning tree is a two-step process.

- [Configuring Interfaces for Layer 2 Mode](#)
- [Enabling Spanning Tree Protocol Globally](#)

Related Configuration Tasks

- [Adding an Interface to the Spanning Tree Group](#)
- [Removing an Interface from the Spanning Tree Group](#)
- [Modifying Global Parameters](#)
- [Modifying Interface STP Parameters](#)
- [Enabling PortFast](#)
- [Prevent Network Disruptions with BPDU Guard](#)
- [STP Root Guard](#)

- [SNMP Traps for Root Elections and Topology Changes](#)

Important Points to Remember

- STP is disabled by default.
- The Dell Networking operating system (OS) supports only one spanning tree instance (0). For multiple instances, enable the multiple spanning tree protocol (MSTP) or per-VLAN spanning tree plus (PVST+). You may only enable one flavor of spanning tree at any one time.
- All ports in virtual local area networks (VLANs) and all enabled interfaces in Layer 2 mode are automatically added to the spanning tree topology at the time you enable the protocol.
- To add interfaces to the spanning tree topology after you enable STP, enable the port and configure it for Layer 2 using the `switchport` command.
- The IEEE Standard 802.1D allows 8 bits for port ID and 8 bits for priority. The 8 bits for port ID provide port IDs for 256 ports.

Configuring Interfaces for Layer 2 Mode

All interfaces on all switches that participate in spanning tree must be in Layer 2 mode and enabled.

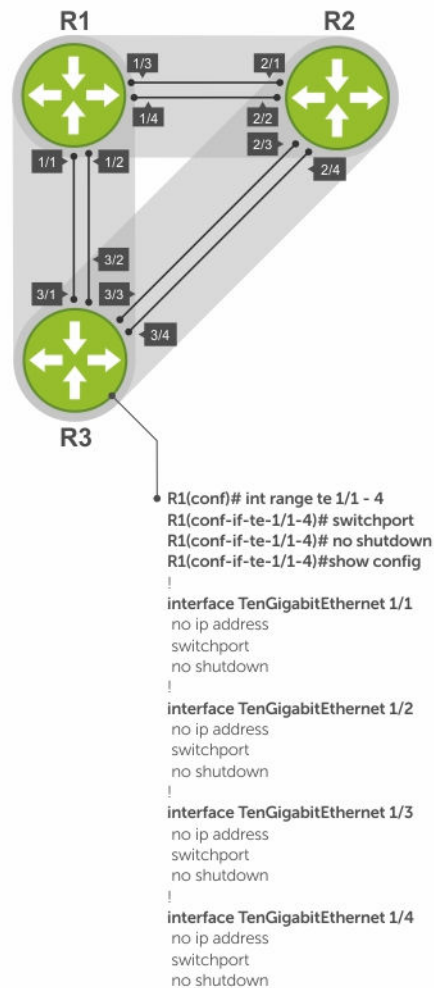


Figure 121. Example of Configuring Interfaces for Layer 2 Mode

To configure and enable the interfaces for Layer 2, use the following command.

1. If the interface has been assigned an IP address, remove it.
INTERFACE mode

no ip address
2. Place the interface in Layer 2 mode.
INTERFACE

switchport

3. Enable the interface.

INTERFACE mode

```
no shutdown
```

Example of the `show config` Command

To verify that an interface is in Layer 2 mode and enabled, use the `show config` command from INTERFACE mode.

```
Dell(conf-if-te-1/1)#show config
!
interface TenGigabitEthernet 1/1
  no ip address
  switchport
  no shutdown
Dell(conf-if-te-1/1)#
```

Enabling Spanning Tree Protocol Globally

Enable the spanning tree protocol globally; it is not enabled by default.

When you enable STP, all physical, VLAN, and port-channel interfaces that are enabled and in Layer 2 mode are automatically part of the Spanning Tree topology.

- Only one path from any bridge to any other bridge participating in STP is enabled.
- Bridges block a redundant path by disabling one of the link ports.




Figure 122. Spanning Tree Enabled Globally

To enable STP globally, use the following commands.

1. Enter PROTOCOL SPANNING TREE mode.

CONFIGURATION mode

```
protocol spanning-tree 0
```

2. Enable STP.

PROTOCOL SPANNING TREE mode

```
no disable
```

Example of Verifying Spanning Tree is Enabled

Example of Viewing Spanning Tree Configuration

Example of Verifying a Port Participates in Spanning Tree

To disable STP globally for all Layer 2 interfaces, use the `disable` command from PROTOCOL SPANNING TREE mode.

To verify that STP is enabled, use the `show config` command from PROTOCOL SPANNING TREE mode.

```
Dell(conf)#protocol spanning-tree 0
Dell(conf-span)#show config
!
protocol spanning-tree 0
  no disable
Dell#
```

To view the spanning tree configuration and the interfaces that are participating in STP, use the `show spanning-tree 0` command from EXEC privilege mode. If a physical interface is part of a port channel, only the port channel is listed in the command output.

```
R2#show spanning-tree 0
Executing IEEE compatible Spanning Tree Protocol
  Bridge Identifier has priority 32768, address 0001.e826.ddb7
  Configured hello time 2, max age 20, forward delay 15
  Current root has priority 32768, address 0001.e80d.2462
  Root Port is 289 (GigabitEthernet 2/1), cost of root path is 4
  Topology change flag not set, detected flag not set
  Number of topology changes 3 last change occurred 0:16:11 ago
    from GigabitEthernet 2/3
  Timers: hold 1, topology change 35
          hello 2, max age 20, forward delay 15
  Times: hello 0, topology change 0, notification 0, aging Normal

Port 289 (GigabitEthernet 2/1) is Forwarding
  Port path cost 4, Port priority 8, Port Identifier 8.289
  Designated root has priority 32768, address 0001.e80d.2462
  Designated bridge has priority 32768, address 0001.e80d.2462
  Designated port id is 8.496, designated path cost 0
  Timers: message age 1, forward delay 0, hold 0
  Number of transitions to forwarding state 1
  BPDU: sent 21, received 486
  The port is not in the portfast mode

Port 290 (GigabitEthernet 2/2) is Blocking
  Port path cost 4, Port priority 8, Port Identifier 8.290
--More--
  Timers: message age 1, forward delay 0, hold 0
  Number of transitions to forwarding state 1
  BPDU: sent 21, received 486
  The port is not in the portfast mode
```

To confirm that a port is participating in Spanning Tree, use the `show spanning-tree 0 brief` command from EXEC privilege mode.

```
Dell#show spanning-tree 0 brief
Executing IEEE compatible Spanning Tree Protocol
  Root ID Priority 32768, Address 0001.e80d.2462
  We are the root of the spanning tree
  Root Bridge hello time 2, max age 20, forward delay 15
  Bridge ID Priority 32768, Address 0001.e80d.2462
  Configured hello time 2, max age 20, forward delay 15
```

Interface Name	PortID	Prio	Cost	Sts	Cost	Designated Bridge ID	PortID
Tengig 1/1	8.496	8	4	DIS	0	32768 0001.e80d.2462	8.496
Tengig 1/2	8.497	8	4	DIS	0	32768 0001.e80d.2462	8.497
Tengig 1/3	8.513	8	4	FWD	0	32768 0001.e80d.2462	8.513
Tengig 1/4	8.514	8	4	FWD	0	32768 0001.e80d.2462	8.514

```
Dell#
```

Adding an Interface to the Spanning Tree Group

To add a Layer 2 interface to the spanning tree topology, use the following command.

- Enable spanning tree on a Layer 2 interface.

```
INTERFACE mode  
  
spanning-tree 0
```

Removing an Interface from the Spanning Tree Group


To remove a Layer 2 interface from the spanning tree topology, use the following command.

- Disable spanning tree on a Layer 2 interface.

```
INTERFACE mode  
  
no spanning-tree 0
```

Modifying Global Parameters

You can modify the spanning tree parameters. The root bridge sets the values for forward-delay, hello-time, and max-age and overwrites the values set on other bridges participating in STP.

 **NOTE:** Dell Networking recommends that only experienced network administrators change the spanning tree parameters. Poorly planned modification of the spanning tree parameters can negatively affect network performance.

The following table displays the default values for STP.

Table 68. STP Default Values

STP Parameters	Default Value
Forward Delay	15 seconds
Hello Time	2 seconds
Max Age	20 seconds
Port Cost	
• 40-Gigabit Ethernet interfaces	• 1
• 10-Gigabit Ethernet interfaces	• 2
• Port Channel with 40-Gigabit Ethernet interfaces	• 1
• Port Channel with 10-Gigabit Ethernet interfaces	• 1
Port Priority	8

- Change the `forward-delay` parameter (the wait time before the interface enters the Forwarding state).

PROTOCOL SPANNING TREE mode

```
forward-delay seconds
```


The range is from 4 to 30.

The default is **15 seconds**.

- Change the `hello-time` parameter (the BPDU transmission interval).

PROTOCOL SPANNING TREE mode

```
hello-time seconds
```

 **NOTE:** With large configurations (especially those with more ports) Dell Networking recommends increasing the hello-time.

The range is from 1 to 10.

the default is **2 seconds**.

- Change the `max-age` parameter (the refresh interval for configuration information that is generated by recomputing the spanning tree topology).

PROTOCOL SPANNING TREE mode

```
max-age seconds
```

The range is from 6 to 40.

The default is **20 seconds**.

To view the current values for global parameters, use the `show spanning-tree 0` command from EXEC privilege mode. Refer to the second example in [Enabling Spanning Tree Protocol Globally](#).

Modifying Interface STP Parameters

You can set the port cost and port priority values of interfaces in Layer 2 mode.

- **Port cost** — a value that is based on the interface type. The greater the port cost, the less likely the port is selected to be a forwarding port.
- **Port priority** — influences the likelihood that a port is selected to be a forwarding port in case that several ports have the same port cost.

To change the port cost or priority of an interface, use the following commands.

- Change the port cost of an interface.

INTERFACE mode

```
spanning-tree 0 cost cost
```

The range is from 0 to 65535.

The default values are listed in [Modifying Global Parameters](#).

- Change the port priority of an interface.

INTERFACE mode


```
spanning-tree 0 priority priority-value
```

The range is from 0 to 15.

The default is **8**.

To view the current values for interface parameters, use the `show spanning-tree 0` command from EXEC privilege mode. Refer to the second example in [Enabling Spanning Tree Protocol Globally](#).

Enabling PortFast

The PortFast feature enables interfaces to begin forwarding traffic approximately 30 seconds sooner. Interfaces forward frames by default until they receive a BPDU that indicates that they should behave otherwise; they do not go through the Learning and Listening states. The `bpduguard shutdown-on-violation` option causes the interface hardware to be shut down when it receives a BPDU. When you only implement `bpduguard`, although the interface is placed in an Error Disabled state when receiving the BPDU, the physical interface remains up and spanning-tree drops packets in the hardware after a BPDU violation. BPDUs are dropped in the software after receiving the BPDU violation.

PDU Filtering enabled on an interface stops sending and receiving BPDUs on the port fast enabled ports. When you enable BPDU guard and BPDU filter on the port, then BPDU filter takes the highest precedence. By default BPDU filtering on an interface is disabled.

 **CAUTION: Enable PortFast only on links connecting to an end station. PortFast can cause loops if it is enabled on an interface connected to a network.**

To enable PortFast on an interface, use the following command.

- Enable PortFast on an interface.

```
INTERFACE mode
```

```
spanning-tree stp-id portfast [bpduguard | [shutdown-on-violation] |  
bpdupfilter]
```

Example of Verifying PortFast is Enabled on an Interface

To verify that PortFast is enabled on a port, use the `show spanning-tree` command from EXEC Privilege mode or the `show config` command from INTERFACE mode. Dell Networking recommends using the `show config` command.

```
Dell#(conf-if-te-1/1)#show conf  
!  
interface Tengigabitethernet 1/1  
  no ip address  
  switchport  
  spanning-tree 0 portfast  
  no shutdown  
Dell#(conf-if-te-1/1)#
```

Prevent Network Disruptions with BPDU Guard


Configure the Portfast (and Edgeport, in the case of RSTP, PVST+, and MSTP) feature on ports that connect to end stations. End stations do not generate BPDUs, so ports configured with Portfast/ Edgeport (edgeports) do not expect to receive BPDUs.

If an edgeport does receive a BPDU, it likely means that it is connected to another part of the network, which can negatively affect the STP topology. The BPDU Guard feature blocks an edgeport after receiving a BPDU to prevent network disruptions, and the system displays the following message.

```
3w3d0h: %RPM0-P:RP2 %SPANMGR-5-BPDU_GUARD_RX_ERROR: Received Spanning Tree BPDU
on
BPDU guard port. Disable GigabitEthernet 3/41.
```

Enable BPDU Guard using the `bpduguard` option when enabling PortFast or EdgePort. The `bpduguard shutdown-on-violation` option causes the interface hardware to be shut down when it receives a BPDU. Otherwise, although the interface is placed in an Error Disabled state when receiving the BPDU, the physical interface remains up and spanning-tree will only drop packets after a BPDU violation.

The following example shows a scenario in which an edgeport might unintentionally receive a BPDU. The port on the Dell Networking system is configured with Portfast. If the switch is connected to the hub, the BPDUs that the switch generates might trigger an undesirable topology change. If you enable BPDU Guard, when the edge port receives the BPDU, the BPDU is dropped, the port is blocked, and a console message is generated.

 **NOTE:** Unless you enable the `shutdown-on-violation` option, spanning-tree only drops packets after a BPDU violation; the physical interface remains up.

Dell Networking OS Behavior: Regarding `bpduguard shutdown-on-violation` behavior:

- If the interface to be shut down is a port channel, all the member ports are disabled in the hardware.
- When you add a physical port to a port channel already in the Error Disable state, the new member port is also disabled in the hardware.
- When you remove a physical port from a port channel in the Error Disable state, the Error Disabled state is cleared on this physical port (the physical port is enabled in the hardware).
- You can clear the Error Disabled state with any of the following methods:
 - Perform a `shutdown` command on the interface.
 - Disable the `shutdown-on-violation` command on the interface (the `no spanning-tree stp-id portfast [bpduguard | [shutdown-on-violation]]` command).
 - Disable spanning tree on the interface (the `no spanning-tree` command in INTERFACE mode).
 - Disabling global spanning tree (the `no spanning-tree` in CONFIGURATION mode).

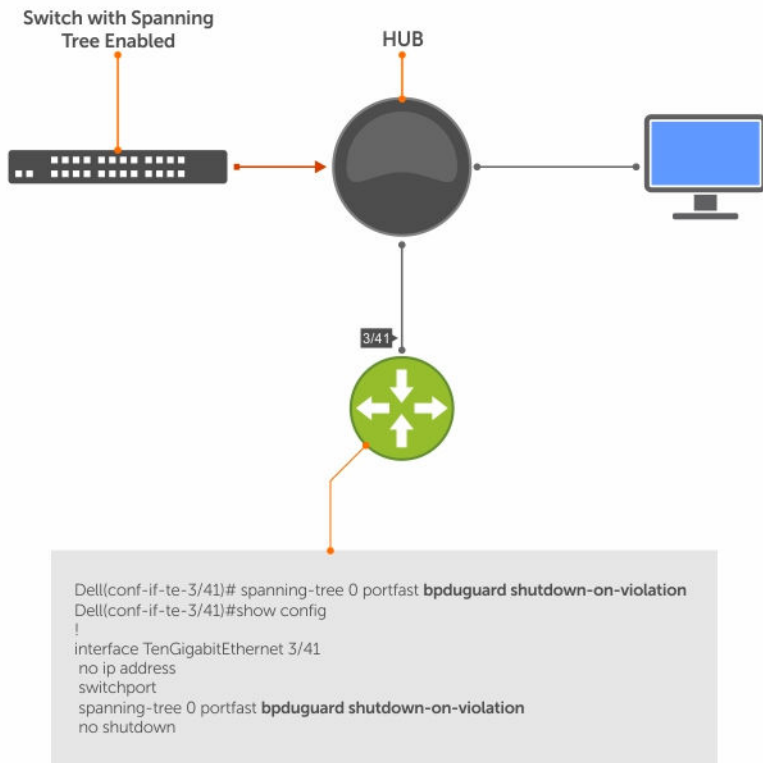


Figure 123. Enabling BPDUGuard

Dell Networking OS Behavior: BPDUGuard and BPDUGuard filtering (refer to [Removing an Interface from the Spanning Tree Group](#)) both block BPDUs, but are two separate features.

BPDUGuard is used on edgeports and blocks all traffic on edgeport if it receives a BPDU.

Example of Blocked BPDUs

```

Dell#show spanning-tree 0 brief
  Executing IEEE compatible Spanning Tree Protocol
  Root ID Priority 32768, Address 0001.e88a.fdb3 Cost 1
  Root Port 2 (Port-channel 1)
  Root Bridge hello time 2, max age 20, forward delay 15
  Bridge ID Priority 32768, Address 001e.c9f1.00cf
  Configured hello time 2, max age 20, forward delay 15
  Bpdu filter disabled globally
Interface          Designated
Name      PortID Prio Cost Sts Cost Bridge ID          PortID
-----
Po 1      8.2   8    1   FWD  0   32768 0001.e88a.fdb3    8.2
Te 3/20   8.317 8    4   EDS  1   32768 001e.c9f1.00cf    8.317
Te 4/20   8.373 8    4   FWD  1   32768 001e.c9f1.00cf    8.373
Te 4/21   8.374 8    4   FWD  1   32768 001e.c9f1.00cf    8.374
Dell#show ip int br ten 3/20
Interface IP-Address OK Method Status Protocol
TenGigabitEthernet 3/20 unassigned YES None up up
Dell#

```

Global BPDU Filtering

When BPDU Filtering is enabled globally, it stops transmitting BPDUs on the operational port fast enabled ports by default.

When it receives BPDUs, it automatically participates in the spanning tree. By default global bpdu filtering is disabled.

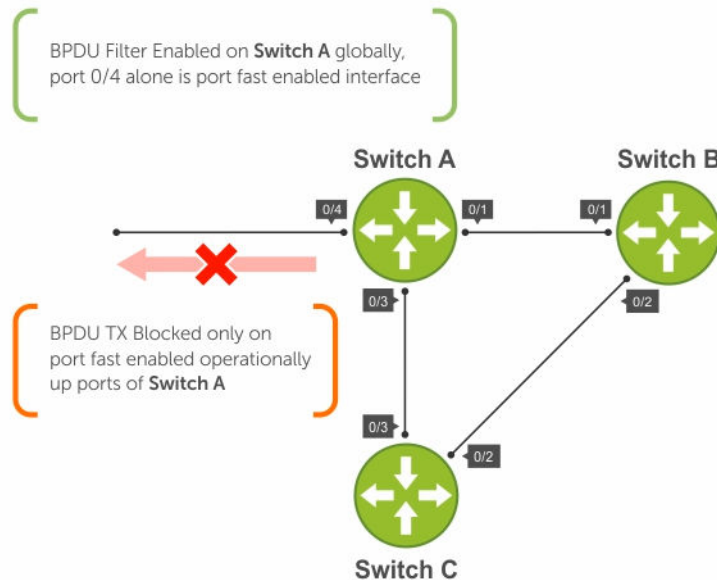


Figure 124. BPDU Filtering Enabled Globally

Interface BPDU Filtering

When BPDU Filtering is enabled on an interface, it should stop sending and receiving BPDUs on the port fast enabled ports.

When BPDU guard and BPDU filter is enabled on the port, then BPDU filter takes the highest precedence. By default bpdu filtering on an interface is disabled.

Add your section content here.

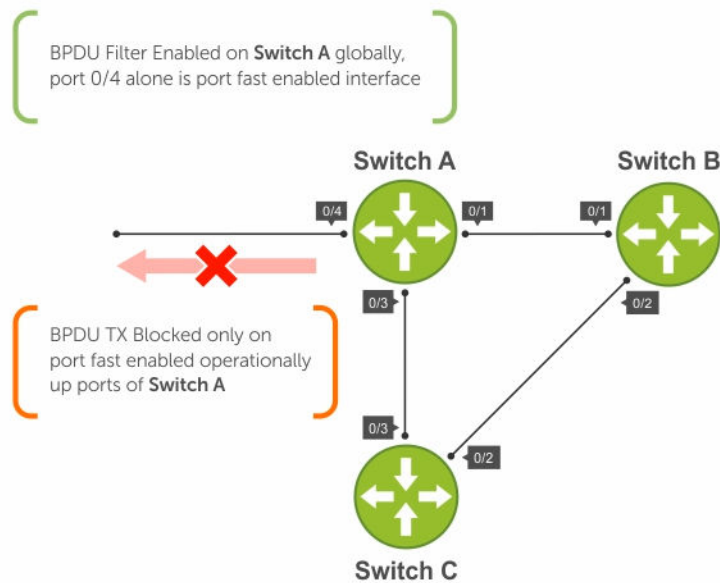


Figure 125. BPDUs Filtering Enabled Globally

Selecting STP Root

The STP determines the root bridge, but you can assign one bridge a lower priority to increase the likelihood that it becomes the root bridge. You can also specify that a bridge is the root or the secondary root.

To change the bridge priority or specify that a bridge is the root or secondary root, use the following command.

- Assign a number as the bridge priority or designate it as the root or secondary root.
PROTOCOL SPANNING TREE mode

```
bridge-priority {priority-value | primary | secondary}
```

- *priority-value*: the range is from 0 to 65535. The lower the number assigned, the more likely this bridge becomes the root bridge.

The primary option specifies a bridge priority of 8192.

The secondary option specifies a bridge priority of 16384.

The default is **32768**.

Example of Viewing STP Root Information

To view only the root information, use the `show spanning-tree root` command from EXEC privilege mode.

```
Dell#show spanning-tree 0 root
Root ID Priority 32768, Address 0001.e80d.2462
We are the root of the spanning tree
```

```
Root Bridge hello time 2, max age 20, forward delay 15
Dell#
```

STP Root Guard

Use the STP root guard feature in a Layer 2 network to avoid bridging loops.

In STP, the switch in the network with the lowest priority (as determined by STP or set with the `bridge-priority` command) is selected as the root bridge. If two switches have the same priority, the switch with the lower MAC address is selected as the root. All other switches in the network use the root bridge as the reference used to calculate the shortest forwarding path.

Because any switch in an STP network with a lower priority can become the root bridge, the forwarding topology may not be stable. The location of the root bridge can change, resulting in unpredictable network behavior. The STP root guard feature ensures that the position of the root bridge does not change.

Root Guard Scenario

For example, as shown in the following illustration (STP topology 1, upper left) Switch A is the root bridge in the network core. Switch C functions as an access switch connected to an external device. The link between Switch C and Switch B is in a Blocking state. The flow of STP BPDUs is shown in the illustration.

In STP topology 2 (shown in the upper right), STP is enabled on device D on which a software bridge application is started to connect to the network. Because the priority of the bridge in device D is lower than the root bridge in Switch A, device D is elected as root, causing the link between Switches A and B to enter a Blocking state. Network traffic then begins to flow in the directions indicated by the BPDUs in the topology. If the links between Switches C and A or Switches C and B cannot handle the increased traffic flow, frames may be dropped.

In STP topology 3 (shown in the lower middle), if you have enabled the root guard feature on the STP port on Switch C that connects to device D, and device D sends a superior BPDUs that would trigger the election of device D as the new root bridge, the BPDUs are ignored and the port on Switch C transitions from a forwarding to a root-inconsistent state (shown by the green X icon). As a result, Switch A becomes the root bridge.

All incoming and outgoing traffic is blocked on an STP port in a Root-Inconsistent state. After the timeout period, the Switch C port automatically transitions to a Forwarding state as soon as device D stops sending BPDUs that advertise a lower priority.

If you enable a root guard on all STP ports on the links where the root bridge should not appear, you can ensure a stable STP network topology and avoid bridging loops.




Figure 126. STP Root Guard Prevents Bridging Loops

Configuring Root Guard

Enable STP root guard on a per-port or per-port-channel basis.

Dell Networking OS Behavior: The following conditions apply to a port enabled with STP root guard:

- Root guard is supported on any STP-enabled port or port-channel interface except when used as a stacking port.
- Root guard is supported on a port in any Spanning Tree mode:
 - [Spanning Tree Protocol \(STP\)](#)
 - [Rapid Spanning Tree Protocol \(RSTP\)](#)
 - [Multicast Source Discovery Protocol \(MSDP\)](#)
 - [Per-VLAN Spanning Tree Plus \(PVST+\)](#)
- When enabled on a port, root guard applies to all VLANs configured on the port.
- When used in an MSTP network, if root guard blocks a boundary port in the CIST, the port is also blocked in all other MST instances.

To enable the root guard on an STP-enabled port or port-channel interface in instance 0, use the following command.

- Enable root guard on a port or port-channel interface.
INTERFACE mode or INTERFACE PORT-CHANNEL mode

```
spanning-tree {0 | mstp | rstp | pvst} rootguard
```

- 0: enables root guard on an STP-enabled port assigned to instance 0.
- mstp: enables root guard on an MSTP-enabled port.
- rstp: enables root guard on an RSTP-enabled port.
- pvst: enables root guard on a PVST-enabled port.

To disable STP root guard on a port or port-channel interface, use the `no spanning-tree 0 rootguard` command in an interface configuration mode.

To verify the STP root guard configuration on a port or port-channel interface, use the `show spanning-tree 0 guard [interface interface]` command in a global configuration mode.

SNMP Traps for Root Elections and Topology Changes

To enable SNMP traps, use the following commands.

- Enable SNMP traps for STP state changes.
`snmp-server enable traps stp`
- Enable SNMP traps for RSTP, MSTP, and PVST+ collectively.
`snmp-server enable traps xstp`

Displaying STP Guard Configuration

To display the STP guard configuration, use the following command. The following example shows an STP network (instance 0) in which:

- Root guard is enabled on a port that is in a root-inconsistent state.
- Loop guard is enabled on a port that is in a listening state.
- BPDU guard is enabled on a port that is shut down (Error Disabled state) after receiving a BPDU.
- Verify the STP guard configured on port or port-channel interfaces.
`show spanning-tree 0 guard [interface interface]`

Example of Viewing STP Guard Configuration

```
Dell#show spanning-tree 0 guard
Interface
Name      Instance Sts          Guard type    Bpdu Filter
-----
Tengig 0/1  0          INCON(Root)   Rootguard    No
Tengig 0/2  0          LIS           Loopguard    No
Tengig 0/3  0          EDS (Shut)    Bpduguard    No
Dell#
```


System Time and Date

System time and date settings and the network time protocol (NTP) are supported on the MXL switch platform.

You can set system times and dates and maintained through the NTP. They are also set through the Dell Networking operating system (OS) command line interfaces (CLIs) and hardware settings.

Network Time Protocol

The network time protocol (NTP) synchronizes timekeeping among a set of distributed time servers and clients.

The protocol also coordinates time distribution in a large, diverse network with various interfaces. In NTP, servers maintain the time and NTP clients synchronize with a time-serving host. NTP clients choose from among several NTP servers to determine which offers the best available source of time and the most reliable transmission of information.

NTP is a fault-tolerant protocol that automatically selects the best of several available time sources to synchronize to. You can combine multiple candidates to minimize the accumulated error. Temporarily or permanently insane time sources are detected and avoided.

Dell Networking recommends configuring NTP for the most accurate time. In the Dell Networking OS, you can configure other time sources (the hardware clock and the software clock).

NTP is designed to produce three products: clock offset, roundtrip delay, and dispersion, all of which are relative to a selected reference clock.

- **Clock offset** — represents the amount to adjust the local clock to bring it into correspondence with the reference clock.
- **Roundtrip delay** — provides the capability to launch a message to arrive at the reference clock at a specified time.
- **Dispersion** — represents the maximum error of the local clock relative to the reference clock.

Because most host time servers synchronize via another peer time server, there are two components in each of these three products, those determined by the peer relative to the primary reference source of standard time and those measured by the host relative to the peer.

In order to facilitate error control and management of the subnet itself, each of these components is maintained separately in the protocol. They provide not only precision measurements of offset and delay, but also definitive maximum error bounds, so that the user interface can determine not only the time, but the quality of the time as well.

In what may be the most common client/server model, a client sends an NTP message to one or more servers and processes the replies as received. The server interchanges addresses and ports, overwrites

certain fields in the message, recalculates the checksum and returns the message immediately. Information included in the NTP message allows the client to determine the server time regarding local time and adjust the local clock accordingly. In addition, the message includes information to calculate the expected timekeeping accuracy and reliability, as well as select the best from possibly several servers.

Following conventions established by the telephone industry [BEL86], the accuracy of each server is defined by a number called the stratum, with the topmost level (primary servers) assigned as one and each level downwards (secondary servers) in the hierarchy assigned as one greater than the preceding level.

The Dell Networking OS synchronizes with a time-serving host to get the correct time. You can set the system to poll specific NTP time-serving hosts for the current time. From those time-serving hosts, the system chooses one NTP host with which to synchronize and serve as a client to the NTP host. As soon as a host-client relationship is established, the networking device propagates the time information throughout its local network.

Protocol Overview

The NTP messages to one or more servers and processes the replies as received. The server interchanges addresses and ports, fills in or overwrites certain fields in the message, recalculates the checksum, and returns it immediately.

Information included in the NTP message allows each client/server peer to determine the timekeeping characteristics of its other peers, including the expected accuracies of their clocks. Using this information, each peer is able to select the best time from possibly several other clocks, update the local clock, and estimate its accuracy.

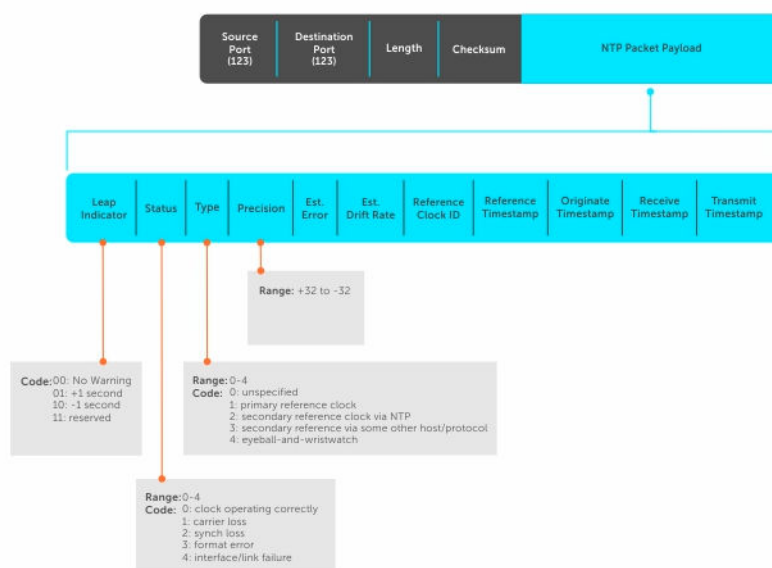


Figure 127. NTP Fields

Implementation Information

Dell Networking systems can only be an NTP client.

Configure the Network Time Protocol

Configuring NTP is a one-step process.

- [Enabling NTP](#)

Related Configuration Tasks

- [Configuring NTP Broadcasts](#)
- [Disabling NTP on an Interface](#)
- [Configuring a Source IP Address for NTP Packets](#)

Enabling NTP

NTP is disabled by default.

To enable NTP, specify an NTP server to which the Dell Networking system synchronizes. To specify multiple servers, enter the command multiple times. You may specify an unlimited number of servers at the expense of CPU resources.

- Specify the NTP server to which the Dell Networking system synchronizes.
CONFIGURATION mode

```
ntp server ip-address
```

Viewing System Clock State Relative to NTP

Example of Viewing Calculated NTP Synchronization Variables

To display the system clock state with respect to NTP, use the `show ntp status` command from EXEC Privilege mode.

```
Dell(conf)#do show ntp status
Clock is synchronized, stratum 2, reference is 192.168.1.1
frequency is -369.623 ppm, stability is 53.319 ppm, precision is 4294967279
reference time is CD63BCC2.0CBBD000 (16:54:26.049 UTC Thu Mar 12 2012)
clock offset is 997.529984 msec, root delay is 0.00098 sec
root dispersion is 10.04271 sec, peer dispersion is 10032.715 msec
peer mode is client
```

To display the calculated NTP synchronization variables received from the server that the system uses to synchronize its clock, use the `show ntp associations` command from EXEC Privilege mode.

```
Dell(conf)#do show ntp associations
remote      ref clock st when poll reach delay offset disp
=====
#192.168.1.1 .LOCL.      1 16 16 76          0.98 -2.470 879.23
* master (syncd), # master (unsyncd), + selected, - candidate
```

Configuring NTP Broadcasts

With the Dell Networking OS, you can receive broadcasts of time information. You can set interfaces within the system to receive NTP information through broadcast. To configure an interface to receive NTP broadcasts, use the following commands.

- Set the interface to receive NTP packets.

```
INTERFACE mode

ntp broadcast client
```

Example of Configuring NTP Broadcasts

```
2w1d11h : NTP: Maximum Slew:-0.000470, Remainder = -0.496884
```

Disabling NTP on an Interface

By default, NTP is enabled on all active interfaces. If you disable NTP on an interface, the system drops any NTP packets sent to that interface.

To disable NTP on an interface, use the following command.

- Disable NTP on the interface.

```
INTERFACE mode

ntp disable
```

To view whether NTP is configured on the interface, use the `show config` command in INTERFACE mode. If `ntp disable` is not listed in the `show config` command output, NTP is enabled. (The `show config` command displays only non-default configuration information.)

Configuring a Source IP Address for NTP Packets

By default, the source address of NTP packets is the IP address of the interface used to reach the network.

You can configure one interface's IP address include in all NTP packets.

To configure an IP address as the source address of NTP packets, use the following command.

- Configure a source IP address for NTP packets.

```
CONFIGURATION mode

ntp source interface
```

Enter the following keywords and slot/port or number information:

- For a Loopback interface, enter the keyword `loopback` then a number between 0 and 16383.
- For a port channel interface, enter the keyword `port-channel` then a number from 1 to 128.
- For a 10-Gigabit Ethernet interface, enter the keyword `TenGigabitEthernet` then the slot/port information.
- For a VLAN interface, enter the keyword `vlan` then a number from 1 to 4094.

- For a 40-Gigabit Ethernet interface, enter the keyword `fortyGigE` then the slot/port information.

To view the configuration, use the `show running-config ntp` command in EXEC privilege mode (refer to the example in [Configuring NTP Authentication](#)).

Configuring NTP Authentication

NTP authentication and the corresponding trusted key provide a reliable means of exchanging NTP packets with trusted time sources.

NTP authentication begins when the first NTP packet is created following the configuration of keys. NTP authentication in the Dell Networking OS uses the message digest 5 (MD5) algorithm and the key is embedded in the synchronization packet that is sent to an NTP time source.

To configure NTP authentication, use the following commands.

1. Enable NTP authentication.

CONFIGURATION mode

```
ntp authenticate
```

2. Set an authentication key.

CONFIGURATION mode

```
ntp authentication-key number md5 key
```

Configure the following parameters:

- *number*: the range is from 1 to 4294967295. This *number* must be the same as the *number* in the `ntp trusted-key` command.
- *key*: enter a text string. This text string is encrypted.

3. Define a trusted key.

CONFIGURATION mode

```
ntp trusted-key number
```

Configure a number from 1 to 4294967295.

The *number* must be the same as the *number* used in the `ntp authentication-key` command.

4. Configure an NTP server.

CONFIGURATION mode

```
ntp server [vrf] <vrf-name> {hostname | ipv4-address | ipv6-address} [ key  
keyid] [prefer] [version number]
```

Configure the IP address of a server and the following optional parameters:

- *vrf-name* : Enter the name of the VRF through which the NTP server is reachable.
- *hostname* : Enter the keyword `hostname` to see the IP address or host name of the remote device.
- *ipv4-address* : Enter an IPv4 address in dotted decimal format (A.B.C.D).
- *ipv6-address* : Enter an IPv6 address in the format 0000:0000:0000:0000:0000:0000:0000:0000. Elision of zeros is supported.

- `key keyid`: Configure a text string as the key exchanged between the NTP server and the client.
- `prefer`: Enter the keyword `prefer` to set this NTP server as the preferred server.
- `version number`: Enter a number as the NTP version. The range is from 1 to 4.

5. Configure the switch as NTP master .

CONFIGURATION mode

```
ntp master <stratum>
```

To configure the switch as NTP Server use the `ntp master<stratum>` command. `stratum` number identifies the NTP Server's hierarchy.

Example of Viewing NTP Configuration

Configuring an NTP Server

To view the NTP configuration, use the `show running-config ntp` command in EXEC privilege mode. The following example shows an encrypted authentication key (in bold). All keys are encrypted.

```
Dell#show running ntp
!
ntp authenticate
ntp authentication-key 345 md5 5A60910F3D211F02
ntp server 11.1.1.1 version 3
ntp trusted-key 345
Dell#
```

```
Dell(conf)#1w6d23h : NTP: xmit packet to 192.168.1.1:
 leap 0, mode 3, version 3, stratum 2, ppoll 1024
 rtdel 0219 (8.193970), rtdsp AF928 (10973.266602), refid C0A80101
(192.168.1.1)
 ref CD7F4F63.6BE8F000 (14:51:15.421 UTC Thu Apr 2 2012)
 org CD7F4F63.68000000 (14:51:15.406 UTC Thu Apr 2 2012)
 rec CD7F4F63.6BE8F000 (14:51:15.421 UTC Thu Apr 2 2012)
 xmt CD7F5368.D0535000 (15:8:24.813 UTC Thu Apr 2 2012)
1w6d23h : NTP: rcv packet from 192.168.1.1
 leap 0, mode 4, version 3, stratum 1, ppoll 1024
 rtdel 0000 (0.000000), rtdsp AF587 (10959.090820), refid 4C4F434C
(76.79.67.76)
 ref CD7E14FD.43F7CED9 (16:29:49.265 UTC Wed Apr 1 2012)
 org CD7F5368.D0535000 (15:8:24.813 UTC Thu Apr 2 2012)
 rec CD7F5368.D0000000 (15:8:24.812 UTC Thu Apr 2 2012)
 xmt CD7F5368.D0000000 (15:8:24.812 UTC Thu Apr 2 2012)
 inp CD7F5368.D1974000 (15:8:24.818 UTC Thu Apr 2 2012)
```

```
rtdel-root delay
rtdsp - round trip dispersion
refid - reference id
org -
rec - (last?) receive timestamp
xmt - transmit timestamp
```

```
mode - 3 client, 4 server
stratum - 1 primary reference clock, 2 secondary reference clock (via NTP)
version - NTP version 3
leap
```

NOTE:

- **Leap Indicator** (`sys.leap`, `peer.leap`, `pkt.leap`) — This is a two-bit code warning of an impending leap second to be inserted in the NTP time scale. The bits are set before 23:59 on the day of insertion and reset after 00:00 on the following day. This causes the number of seconds (rollover interval) in the day of insertion to be increased or decreased by one. In the case of primary servers, the bits are set by operator intervention, while in the case of secondary servers, the bits are set by the protocol. The two bits, bit 0, and bit 1, respectively, are coded as follows:
- **Poll Interval** — integer indicating the minimum interval between transmitted messages, in seconds as a power of two. For instance, a value of six indicates a minimum interval of 64 seconds.
- **Precision** — integer indicating the precision of the various clocks, in seconds to the nearest power of two. The value must be rounded to the next larger power of two; for instance, a 50 Hz (20 ms) or 60 Hz (16.67ms) power-frequency clock is assigned the value -5 (31.25 ms), while a 1000 Hz (1 ms) crystal-controlled clock is assigned the value -9 (1.95 ms).
- **Root Delay** (`sys.rootdelay`, `peer.rootdelay`, `pkt.rootdelay`) — a signed fixed-point number indicating the total round-trip delay to the primary reference source at the root of the synchronization subnet, in seconds. This variable can take on both positive and negative values, depending on clock precision and skew.
- **Root Dispersion** (`sys.rootdispersion`, `peer.rootdispersion`, `pkt.rootdispersion`) — a signed fixed-point number indicating the maximum error relative to the primary reference source at the root of the synchronization subnet, in seconds. Only positive values greater than zero are possible.
- **Reference Clock Identifier** (`sys.refid`, `peer.refid`, `pkt.refid`) — This is a 32-bit code identifying the particular reference clock. In the case of stratum 0 (unspecified) or stratum 1 (primary reference source), this is a four-octet, left-justified, zero-padded ASCII string, for example: in the case of stratum 2 and greater (secondary reference) this is the four-octet internet address of the peer selected for synchronization.
- **Reference Timestamp** (`sys.reftime`, `peer.reftime`, `pkt.reftime`) — This is the local time, in timestamp format, when the local clock was last updated. If the local clock has never been synchronized, the value is zero.
- **Originate Timestamp**: The departure time on the server of its last NTP message. If the server becomes unreachable, the value is set to zero.
- **Receive Timestamp** — the arrival time on the client of the last NTP message from the server. If the server becomes unreachable, the value is set to zero.
- **Transmit Timestamp** — the departure time on the server of the current NTP message from the sender.
- **Filter dispersion** — the error in calculating the minimum delay from a set of sample data from a peer.

Dell Networking OS Time and Date

You can set the time and date using the Dell Networking OS CLI.

Configuration Task List

The following is a configuration task list for configuring the time and date settings.

- [Setting the Time and Date for the Switch Software Clock](#)
- [Setting the Timezone](#)

- [Setting Daylight Saving Time Once](#)
- [Setting Recurring Daylight Saving Time](#)

Setting the Time and Date for the Switch Software Clock

You can change the order of the `month` and `day` parameters to enter the time and date as *time day month year*. You cannot delete the software clock.

The software clock runs only when the software is up. The clock restarts, based on the hardware clock, when the switch reboots.

To set the software clock, use the following command.

- Set the system software clock to the current time and date.

EXEC Privilege mode

```
clock set time month day year
```

- *time*: enter the time in hours:minutes:seconds. For the hour variable, use the 24-hour format; for example, 17:15:00 is 5:15 pm.
- *month*: enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to *time day month year*.
- *day*: enter the number of the day. The range is from 1 to 31. You can enter the name of a month to change the order of the display to *time day month year*.
- *year*: enter a four-digit number as the year. The range is from 1993 to 2035.

Example of the `clock set` Command

```
Dell#clock set 12:11:00 21 may 2012
Dell#
```

Setting the Timezone

Universal time coordinated (UTC) is the time standard based on the International Atomic Time standard, commonly known as Greenwich Mean time.

When determining system time, include the differentiator between UTC and your local timezone. For example, San Jose, CA is the Pacific Timezone with a UTC offset of -8.

To set the clock timezone, use the following command.

- Set the clock to the appropriate timezone.

CONFIGURATION mode

```
clock timezone timezone-name offset
```

- *timezone-name*: enter the name of the timezone. Do not use spaces.
- *offset*: enter one of the following:
 - * a number from 1 to 23 as the number of hours in addition to UTC for the timezone.
 - * a minus sign (-) then a number from 1 to 23 as the number of hours.

Example of the `clock timezone` Command

```
Dell#conf
Dell(conf)#clock timezone Pacific -8
Dell#
```


Set Daylight Saving Time

The Dell Networking OS supports setting the system to daylight saving time once or on a recurring basis every year.

Setting Daylight Saving Time Once

Set a date (and time zone) on which to convert the switch to daylight saving time on a one-time basis. To set the clock for daylight savings time once, use the following command.

- Set the clock to the appropriate timezone and daylight saving time.

CONFIGURATION mode

```
clock summer-time time-zone date start-month start-day start-year start-time  
end-month end-day end-year end-time [offset]
```

- *time-zone*: enter the three-letter name for the time zone. This name displays in the show clock output.
- *start-month*: enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to *time day month year*.
- *start-day*: enter the number of the day. The range is from 1 to 31. You can enter the name of a month to change the order of the display to *time day month year*.
- *start-year*: enter a four-digit number as the year. The range is from 1993 to 2035.
- *start-time*: enter the time in hours:minutes. For the hour variable, use the 24-hour format; example, 17:15 is 5:15 pm.
- *end-month*: enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to *time day month year*.
- *end-day*: enter the number of the day. The range is from 1 to 31. You can enter the name of a month to change the order of the display to *time day month year*.
- *end-year*: enter a four-digit number as the year. The range is from 1993 to 2035.
- *end-time*: enter the time in hours:minutes. For the hour variable, use the 24-hour format; example, 17:15 is 5:15 pm.
- *offset*: (OPTIONAL) enter the number of minutes to add during the summer-time period. The range is from 1 to 1440. The default is **60 minutes**.

Example of the `clock summer-time` Command

```
Dell(conf)#clock summer-time pacific date Mar 14 2012 00:00 Nov 7 2012 00:00  
Dell(conf)#
```

Setting Recurring Daylight Saving Time

Set a date (and time zone) on which to convert the switch to daylight saving time on a specific day every year.

If you have already set daylight saving for a one-time setting, you can set that date and time as the recurring setting with the `clock summer-time time-zone recurring` command.

To set a recurring daylight saving time, use the following command.

- Set the clock to the appropriate timezone and adjust to daylight saving time every year.

CONFIGURATION mode


`clock summer-time time-zone recurring start-week start-day start-month start-time end-week end-day end-month end-time [offset]`

- *time-zone*: Enter the three-letter name for the time zone. This name displays in the show clock output.
- *start-week*: (OPTIONAL) Enter one of the following as the week that daylight saving begins and then enter values for *start-day* through *end-time*:
 - * *week-number*: Enter a number from 1 to 4 as the number of the week in the month to start daylight saving time.
 - * *first*: Enter the keyword *first* to start daylight saving time in the first week of the month.
 - * *last*: Enter the keyword *last* to start daylight saving time in the last week of the month.
- *start-month*: Enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to *time day month year*.
- *start-day*: Enter the number of the day. The range is from 1 to 31. You can enter the name of a month to change the order of the display to *time day month year*.
- *start-year*: Enter a four-digit number as the year. The range is from 1993 to 2035.
- *start-time*: Enter the time in hours:minutes. For the hour variable, use the 24-hour format; example, 17:15 is 5:15 pm.
- *end-week*: If you entered a start-week, enter the one of the following as the week that daylight saving ends:
 - * *week-number*: Enter a number from 1 to 4 as the number of the week in the month to start daylight saving time.
 - * *first*: Enter the keyword *first* to start daylight saving time in the first week of the month.
 - * *last*: Enter the keyword *last* to start daylight saving time in the last week of the month.
- *end-month*: Enter the name of one of the 12 months in English. You can enter the name of a day to change the order of the display to *time day month year*.
- *end-day*: Enter the number of the day. The range is from 1 to 31. You can enter the name of a month to change the order of the display to *time day month year*.
- *end-year*: Enter a four-digit number as the year. The range is from 1993 to 2035.
- *end-time*: Enter the time in hours:minutes. For the hour variable, use the 24-hour format; example, 17:15 is 5:15 pm.
- *offset*: (OPTIONAL) Enter the number of minutes to add during the summer-time period. The range is from 1 to 1440. The default is **60 minutes**.

Example of the `clock summer-time recurring` Command

Example of Clock Summer-Time Recurring Parameters

```
Dell(conf)#clock summer-time pacific recurring Mar 14 2012 00:00 Nov 7 2012
00:00
Dell(conf)#
```

 **NOTE:** If you enter <CR> after entering the `recurring` command parameter, and you have already set a one-time daylight saving time/date, the system uses that time and date as the recurring setting.

```
Dell(conf)#clock summer-time pacific recurring ?
<1-4>      Week number to start
first      Week number to start
last       Week number to start
<cr>
```

```
Dell(conf)#clock summer-time pacific recurring
Dell(conf)#
```

Tunneling

Tunneling supports RFC 2003, RFC 2473, and 4213. DSCP, hop-limits, flow label values, OSPFv2, and OSPFv3 are also supported. ICMP error relay, PATH MTU transmission, and fragmented packets are not supported.

Configuring a Tunnel

You can configure a tunnel in IPv6 mode, IPv6IP mode, and IPIP mode.

- If the tunnel mode is IPIP or IPv6IP, the tunnel source address and the tunnel destination address must be an IPv4 address.
- If the tunnel mode is IPv6, the tunnel source address and the tunnel destination address must be an IPv6 address.
- If the tunnel mode is IPv6 or IPIP, you can use either an IPv6 address or an IPv4 address for the logical address of the tunnel, but in IPv6IP mode, the logical address must be an IPv6 address.

The following sample configuration shows a tunnel configured in IPv6 mode (carries IPv6 and IPv4 traffic).

```
Dell(conf)#interface tunnel 1
Dell(conf-if-tu-1)#tunnel source 30.1.1.1
Dell(conf-if-tu-1)#tunnel destination 50.1.1.1
Dell(conf-if-tu-1)#tunnel mode ipip
Dell(conf-if-tu-1)#ip address 1.1.1.1/24
Dell(conf-if-tu-1)#ipv6 address 1::1/64
Dell(conf-if-tu-1)#no shutdown
Dell(conf-if-tu-1)#show config
!
interface Tunnel 1
ip address 1.1.1.1/24
ipv6 address 1::1/64
tunnel destination 50.1.1.1
tunnel source 30.1.1.1
tunnel mode ipip
no shutdown
```

The following sample configuration shows a tunnel configured in IPv6IP mode (IPv4 tunnel carries IPv6 traffic only):

```
Dell(conf)#interface tunnel 2
Dell(conf-if-tu-2)#tunnel source 60.1.1.1
Dell(conf-if-tu-2)#tunnel destination 90.1.1.1
Dell(conf-if-tu-2)#tunnel mode ipv6ip
Dell(conf-if-tu-2)#ipv6 address 2::1/64
Dell(conf-if-tu-2)#no shutdown
Dell(conf-if-tu-2)#show config
!
interface Tunnel 2
no ip address
```

```
ipv6 address 2::1/64
tunnel destination 90.1.1.1
tunnel source 60.1.1.1
tunnel mode ipv6ip
no shutdown
```

The following sample configuration shows a tunnel configured in IPIP mode (IPv4 tunnel carries IPv4 and IPv6 traffic):

```
Dell(conf)#interface tunnel 3
Dell(conf-if-tu-3)#tunnel source 5::5
Dell(conf-if-tu-3)#tunnel destination 8::9
Dell(conf-if-tu-3)#tunnel mode ipv6
Dell(conf-if-tu-3)#ip address 3.1.1.1/24
Dell(conf-if-tu-3)#ipv6 address 3::1/64
Dell(conf-if-tu-3)#no shutdown
Dell(conf-if-tu-3)#show config
!
interface Tunnel 3
ip address 3.1.1.1/24
ipv6 address 3::1/64
tunnel destination 8::9
tunnel source 5::5
tunnel mode ipv6
no shutdown
```

Configuring Tunnel keepalive

Configure the tunnel keepalive target, interval and attempts.

- By default the tunnel keepalive is disabled.

The following sample configuration shows tunnel keepalive command:

```
Dell(conf-if-te-0/12)#show config
!
interface TenGigabitEthernet 0/12
ip address 40.1.1.1/24
ipv6 address 500:10::1/64
no shutdown
Dell(conf-if-te-0/12)#
Dell(conf)#interface tunnel 1
Dell(conf-if-tu-1)#ipv6 address 1abd::1/64
Dell(conf-if-tu-1)#ip address 1.1.1.1/24
Dell(conf-if-tu-1)#tunnel source 40.1.1.1
Dell(conf-if-tu-1)#tunnel destination 40.1.1.2
Dell(conf-if-tu-1)#tunnel mode ipip
Dell(conf-if-tu-1)#no shutdown
Dell(conf-if-tu-1)#tunnel keepalive 1.1.1.2 attempts 4 interval 6
Dell(conf-if-tu-1)#show config
!
interface Tunnel 1
ip address 1.1.1.1/24
ipv6 address 1abd::1/64
tunnel destination 40.1.1.2
tunnel source 40.1.1.1
tunnel keepalive 1.1.1.2 attempts 4 interval 6
```

```
tunnel mode ipip
no shutdown
```

Configuring the ip and ipv6 unnumbered

Configuring the tunnel interface is supported on the MXL platform.

You can configure the tunnel in ip unnumbered and ipv6 unnumbered command. To configure the tunnel interface to operate without a unique explicit ip/ ipv6 address, select the interface from which the tunnel will borrow its address.

The following sample configuration shows the IP unnumbered command:

```
Dell(conf-if-te-0/0)#show config
!
interface TenGigabitEthernet 0/0
ip address 20.1.1.1/24
ipv6 address 20:1::1/64
no shutdown
Dell(conf)#interface tunnel 1
Dell(conf-if-tu-1)#ip unnumbered tengigabitethernet 0/0
Dell(conf-if-tu-1)#ipv6 unnumbered tengigabitethernet 0/0
Dell(conf-if-tu-1)#tunnel source 40.1.1.1
Dell(conf-if-tu-1)#tunnel mode ipip decapsulate-any
Dell(conf-if-tu-1)#no shutdown
Dell(conf-if-tu-1)#sho c
!
interface Tunnel 1
ip unnumbered TenGigabitEthernet 0/0
ipv6 unnumbered TenGigabitEthernet 0/0
tunnel source 40.1.1.1
tunnel mode ipip decapsulate-any
no shutdown
Dell(conf-if-tu-1)#
```

Configuring the Tunnel allow-remote

You can configure an IPv4 or IPV6 address or prefix whose tunneled packet will be accepted for decapsulation.

.

- If no allow-remote entries are configured, then tunneled packets from any remote peer address will be accepted.
- Upto eight allow-remote entries can be configured on any particular multipoint receive-only tunnel.

The following sample configuration shows tunnel allow-remote command:

```
Dell(conf)#interface tunnel 1
Dell(conf-if-tu-1)#ipv6 address 1abd::1/64
Dell(conf-if-tu-1)#ip address 1.1.1.1/24
Dell(conf-if-tu-1)#tunnel source 40.1.1.1
Dell(conf-if-tu-1)#tunnel mode ipip decapsulate-any
Dell(conf-if-tu-1)#tunnel allow-remote 40.1.1.2
Dell(conf-if-tu-1)#no shutdown
Dell(conf-if-tu-1)#show config
!
interface Tunnel 1
```

```
ip address 1.1.1.1/24
ipv6 address labd::1/64
tunnel source 40.1.1.1
tunnel allow-remote 40.1.1.2
tunnel mode ipip decapsulate-any
no shutdown
```

Configuring the Tunnel Source Anylocal

You can use the anylocal argument in place of the ip address or interface, but only with multipoint receive-only mode tunnels. The `tunnel source anylocal` command allows the multipoint receive-only tunnel to decapsulate tunnel packets addressed to any IPv4 or IPv6 (depending on the tunnel mode) address configured on the switch that is operationally UP.

The following sample configuration shows the `tunnel source anylocal` command:

```
Dell(conf)#interface tunnel 1
Dell(conf-if-tu-1)#ipv6 address labd::1/64
Dell(conf-if-tu-1)#ip address 1.1.1.1/24
Dell(conf-if-tu-1)#tunnel source anylocal
Dell(conf-if-tu-1)#tunnel mode ipip decapsulate-any
Dell(conf-if-tu-1)#tunnel allow-remote 40.1.1.2
Dell(conf-if-tu-1)#no shutdown
Dell(conf-if-tu-1)#show config
!
interface Tunnel 1
ip address 1.1.1.1/24
ipv6 address labd::1/64
tunnel source anylocal
tunnel allow-remote 40.1.1.2
tunnel mode ipip decapsulate-any
no shutdown
```

Uplink Failure Detection (UFD)

Uplink failure detection (UFD) is supported on the MXL switch platform.

Feature Description

UFD provides detection of the loss of upstream connectivity and, if used with network interface controller (NIC) teaming, automatic recovery from a failed link.

A switch provides upstream connectivity for devices, such as servers. If a switch loses its upstream connectivity, downstream devices also lose their connectivity. However, the devices do not receive a direct indication that upstream connectivity is lost because connectivity to the switch is still operational.

UFD allows a switch to associate downstream interfaces with upstream interfaces. When upstream connectivity fails, the switch disables the downstream links. Failures on the downstream links allow downstream devices to recognize the loss of upstream connectivity.

For example, as shown in the following illustration, Switches S1 and S2 both have upstream connectivity to Router R1 and downstream connectivity to the server. UFD operation is shown in Steps A through C:

- In Step A, the server configuration uses the connection to S1 as the primary path. Network traffic flows from the server to S1 and then upstream to R1.
- In Step B, the upstream link between S1 and R1 fails. The server continues to use the link to S1 for its network traffic, but the traffic is not successfully switched through S1 because the upstream link is down.
- In Step C, UFD on S1 disables the link to the server. The server then stops using the link to S1 and switches to using its link to S2 to send traffic upstream to R1.




Figure 128. Uplink Failure Detection

How Uplink Failure Detection Works

UFD creates an association between upstream and downstream interfaces. The association of uplink and downlink interfaces is called an *uplink-state group*.

An interface in an uplink-state group can be a physical interface or a port-channel (LAG) aggregation of physical interfaces.

An enabled uplink-state group tracks the state of all assigned upstream interfaces. Failure on an upstream interface results in the automatic disabling of downstream interfaces in the uplink-state group. As a

result, downstream devices can execute the protection or recovery procedures they have in place to establish alternate connectivity paths, as shown in the following illustration.

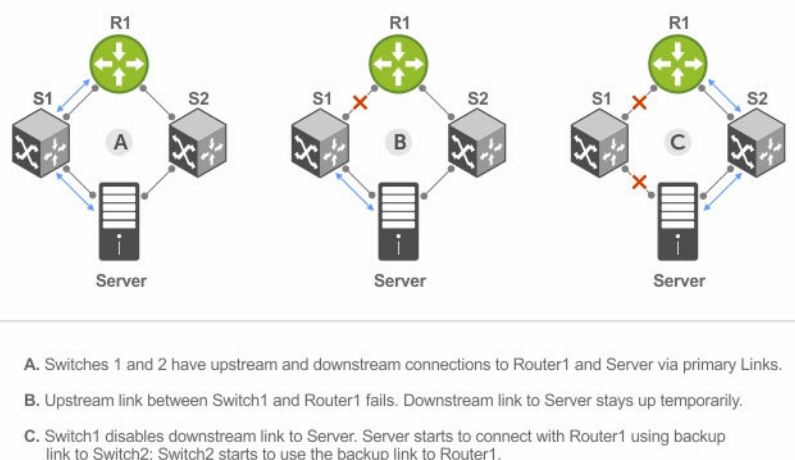


Figure 129. Uplink Failure Detection Example

If only one of the upstream interfaces in an uplink-state group goes down, a specified number of downstream ports associated with the upstream interface are put into a Link-Down state. You can configure this number and is calculated by the ratio of the upstream port bandwidth to the downstream port bandwidth in the same uplink-state group. This calculation ensures that there is no traffic drops due to insufficient bandwidth on the upstream links to the routers/switches.

By default, if all upstream interfaces in an uplink-state group go down, all downstream interfaces in the same uplink-state group are put into a Link-Down state.

Using UFD, you can configure the automatic recovery of downstream ports in an uplink-state group when the link status of an upstream port changes. The tracking of upstream link status does not have a major impact on central processing unit (CPU) usage.

UFD and NIC Teaming

To implement a rapid failover solution, you can use uplink failure detection on a switch with network adapter teaming on a server.

For more information, refer to [NIC Teaming](#).

For example, as shown previously, the switch/ router with UFD detects the uplink failure and automatically disables the associated downstream link port to the server. To continue to transmit traffic upstream, the server with NIC teaming detects the disabled link and automatically switches over to the backup link in order to continue to transmit traffic upstream.

Important Points to Remember

When you configure UFD, the following conditions apply.

- You can configure up to 16 uplink-state groups. By default, no uplink-state groups are created.
 - An uplink-state group is considered to be operationally *up* if it has at least one upstream interface in the Link-Up state.
 - An uplink-state group is considered to be operationally *down* if it has no upstream interfaces in the Link-Up state. No uplink-state tracking is performed when a group is disabled or in an Operationally Down state.
- You can assign physical port or port-channel interfaces to an uplink-state group.
 - You can assign an interface to only one uplink-state group. Configure each interface assigned to an uplink-state group as either an upstream or downstream interface, but not both.
 - You can assign individual member ports of a port channel to the group. An uplink-state group can contain either the member ports of a port channel or the port channel itself, but not both.
 - If you assign a port channel as an upstream interface, the port channel interface enters a Link-Down state when the number of port-channel member interfaces in a Link-Up state drops below the configured `minimum number of members` parameter.
- If one of the upstream interfaces in an uplink-state group goes down, either a user-configurable set of downstream ports or all the downstream ports in the group are put in an Operationally Down state with an UFD Disabled error. The order in which downstream ports are disabled is from the lowest numbered port to the highest.
 - If one of the upstream interfaces in an uplink-state group that was down comes up, the set of UFD-disabled downstream ports (which were previously disabled due to this upstream port going down) is brought up and the UFD Disabled error is cleared.
- If you disable an uplink-state group, the downstream interfaces are not disabled regardless of the state of the upstream interfaces.
 - If an uplink-state group has no upstream interfaces assigned, you cannot disable downstream interfaces when an upstream link goes down.
- To enable the debug messages for events related to a specified uplink-state group or all groups, use the `debug uplink-state-group [group-id]` command, where the `group-id` is from 1 to 16.
 - To turn off debugging event messages, use the `no debug uplink-state-group [group-id]` command.
 - For an example of debug log message, refer to [Clearing a UFD-Disabled Interface](#).

Configuring Uplink Failure Detection

To configure UFD, use the following commands.

1. Create an uplink-state group and enable the tracking of upstream links on the switch/router.

CONFIGURATION mode

```
uplink-state-group group-id
```

- *group-id*: values are from 1 to 16.

To delete an uplink-state group, use the `no uplink-state-group group-id` command.

2. Assign a port or port-channel to the uplink-state group as an upstream or downstream interface.
UPLINK-STATE-GROUP mode

```
{upstream | downstream} interface
```

For interface, enter one of the following interface types:

- 10 Gigabit Ethernet: enter `tengigabitethernet {slot/port | slot/port-range}`
- 40 Gigabit Ethernet: enter `fortygigabitethernet {slot/port | slot/port-range}`
- Port channel: enter `port-channel {1-512 | port-channel-range}`

Where *port-range* and *port-channel-range* specify a range of ports separated by a dash (-) and/or individual ports/port channels in any order; for example:

```
upstream gigabitethernet 1/1-2,5,9,11-12
downstream port-channel 1-3,5
```

- A comma is required to separate each port and port-range entry.

To delete an interface from the group, use the `no {upstream | downstream} interface` command.

3. Configure the number of downstream links in the uplink-state group that will be disabled (Oper Down state) if one upstream link in the group goes down.

UPLINK-STATE-GROUP mode

```
downstream disable links {number | all}
```

- *number*: specifies the number of downstream links to be brought down. The range is from 1 to 1024.
- *all*: brings down all downstream links in the group.

The default is no downstream links are disabled when an upstream link goes down.

To revert to the default setting, use the `no downstream disable links` command.

4. (Optional) Enable auto-recovery so that UFD-disabled downstream ports in the uplink-state group come up when a disabled upstream port in the group comes back up.

UPLINK-STATE-GROUP mode

```
downstream auto-recover
```

The default is auto-recovery of UFD-disabled downstream ports is enabled.

To disable auto-recovery, use the `no downstream auto-recover` command.

5. (Optional) Enters a text description of the uplink-state group.

UPLINK-STATE-GROUP mode

```
description text
```

The maximum length is 80 alphanumeric characters.

6. (Optional) Disables upstream-link tracking without deleting the uplink-state group.

UPLINK-STATE-GROUP mode

```
no enable
```

The default is upstream-link tracking is automatically enabled in an uplink-state group.

To re-enable upstream-link tracking, use the `enable` command.

Clearing a UFD-Disabled Interface

You can manually bring up a downstream interface in an uplink-state group that UFD disabled and is in a UFD-Disabled Error state.

To re-enable one or more disabled downstream interfaces and clear the UFD-Disabled Error state, use the following command.

- Re-enable a downstream interface on the switch/router that is in a UFD-Disabled Error State so that it can send and receive traffic.

EXEC mode

```
clear ufd-disable {interface interface | uplink-state-group group-id}
```

For *interface*, enter one of the following interface types:

- 10 Gigabit Ethernet: enter `tengigabitethernet {slot/port | slot/port-range}`
- 40 Gigabit Ethernet: enter `fortygigabitethernet {slot/port | slot/port-range}`
- Port channel: enter `port-channel {1-512 | port-channel-range}`

- * Where *port-range* and *port-channel-range* specify a range of ports separated by a dash (-) and/or individual ports/port channels in any order; for example:

```
tengigabitethernet 1/1-2,5,9,11-12  
port-channel 1-3,5
```

- * A comma is required to separate each port and port-range entry.

```
clear ufd-disable {interface interface | uplink-state-group group-id):re-  
enables all UFD-disabled downstream interfaces in the group. The range is from 1 to 16.
```

Example of Syslog Messages Before and After Entering the `clear ufd-disable uplink-state-group` Command (S50)

The following example message shows the Syslog messages that display when you clear the UFD-Disabled state from all disabled downstream interfaces in an uplink-state group by using the `clear ufd-disable uplink-state-group group-id` command. All downstream interfaces return to an operationally up state.

```
00:10:12: %STKUNIT0-M:CP %IFMGR-5-ASTATE_DN: Changed interface Admin state to  
down: Te 0/1  
00:10:12: %STKUNIT0-M:CP %IFMGR-5-ASTATE_DN: Changed interface Admin state to  
down: Te 0/2  
00:10:12: %STKUNIT0-M:CP %IFMGR-5-ASTATE_DN: Changed interface Admin state to  
down: Te 0/3  
00:10:12: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down:  
Te 0/1  
00:10:12: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down:  
Te 0/2  
00:10:12: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down:  
Te 0/3  
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed uplink state group state
```

```

to down: Group
3
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Downstream interface set to UFD
error-disabled:
Te 0/4
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Downstream interface set to UFD
error-disabled:
Te 0/5
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Downstream interface set to UFD
error-disabled:
Te 0/6
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down:
Te 0/4
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down:
Te 0/5
00:10:13: %STKUNIT0-M:CP %IFMGR-5-OSTATE_DN: Changed interface state to down:
Te 0/6

Dell(conf-if-range-te-0/1-3)#do clear ufd-disable uplink-state-group 3

00:11:50: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Downstream interface cleared from
UFD
error-disabled: Te 0/4
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Downstream interface cleared from
UFD
error-disabled: Te 0/5
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Downstream interface cleared from
UFD
error-disabled: Te 0/6
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Changed interface state to up: Te
0/4
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Changed interface state to up: Te
0/5
00:11:51: %STKUNIT0-M:CP %IFMGR-5-OSTATE_UP: Changed interface state to up: Te
0/6

```

Displaying Uplink Failure Detection

To display information on the UFD feature, use any of the following commands.

- Display status information on a specified uplink-state group or all groups.

EXEC mode

```
show uplink-state-group [group-id] [detail]
```

- *group-id*: The values are 1 to 16.
- *detail*: displays additional status information on the upstream and downstream interfaces in each group.

- Display the current status of a port or port-channel interface assigned to an uplink-state group.

EXEC mode

```
show interfaces interface
```

interface specifies one of the following interface types:

- 10 Gigabit Ethernet: enter `tengigabitethernet slot/port`.
- 40 Gigabit Ethernet: enter `fortygigabitethernet slot/port`.

- Port channel: enter `port-channel {1-512}`.

If a downstream interface in an uplink-state group is disabled (Oper Down state) by uplink-state tracking because an upstream port is down, the message `error-disabled[UFD]` displays in the output.

- Display the current configuration of all uplink-state groups or a specified group.

EXEC mode or UPLINK-STATE-GROUP mode

(For EXEC mode) `show running-config uplink-state-group [group-id]`

(For UPLINK-STATE-GROUP mode) `show configuration`

- `group-id`: The values are from 1 to 16.

Example of Viewing Uplink State Group Status (S50)

Example of Viewing Interface Status with UFD Information (S50)

Examples of Viewing UFD Output

Dell# **show uplink-state-group**

```
Uplink State Group: 1 Status: Enabled, Up
Uplink State Group: 3 Status: Enabled, Up
Uplink State Group: 5 Status: Enabled, Down
Uplink State Group: 6 Status: Enabled, Up
Uplink State Group: 7 Status: Enabled, Up
Uplink State Group: 16 Status: Disabled, Up
```

Dell# **show uplink-state-group 16**

```
Uplink State Group: 16 Status: Disabled, Up
```

Dell#**show uplink-state-group detail**

```
(Up): Interface up (Dwn): Interface down (Dis): Interface disabled
```

```
Uplink State Group      : 1 Status: Enabled, Up
Upstream Interfaces     :
Downstream Interfaces   :
```

```
Uplink State Group      : 3 Status: Enabled, Up
Upstream Interfaces     : Tengig 0/46(Up) Tengig 0/47(Up)
Downstream Interfaces   : Te 13/0(Up) Te 13/1(Up) Te 13/3(Up) Te 13/5(Up) Te
13/6(Up)
```

```
Uplink State Group      : 5 Status: Enabled, Down
Upstream Interfaces     : Tengig 0/0(Dwn) Tengig 0/3(Dwn) Tengig 0/5(Dwn)
Downstream Interfaces   : Te 13/2(Dis) Te 13/4(Dis) Te 13/11(Dis) Te 13/12(Dis)
Te 13/13(Dis)
Te 13/14(Dis) Te 13/15(Dis)
```

```
Uplink State Group      : 6 Status: Enabled, Up
Upstream Interfaces     :
Downstream Interfaces   :
```

```
Uplink State Group      : 7 Status: Enabled, Up
Upstream Interfaces     :
Downstream Interfaces   :
```

```
Uplink State Group      : 16 Status: Disabled, Up
Upstream Interfaces     : Tengig 0/41(Dwn) Po 8(Dwn)
Downstream Interfaces   : Tengig 0/40(Dwn)
```

Dell#**show interfaces gigabitethernet 7/45**

```
TenGigabitEthernet 7/45 is up, line protocol is down (error-disabled[UFD])
```

```

Hardware is Force10Eth, address is 00:01:e8:32:7a:47
  Current address is 00:01:e8:32:7a:47
Interface index is 280544512
Internet address is not set
MTU 1554 bytes, IP MTU 1500 bytes
LineSpeed 1000 Mbit, Mode auto
Flowcontrol rx off tx off
ARP type: ARPA, ARP Timeout 04:00:00
Last clearing of "show interface" counters 00:25:46
Queueing strategy: fifo
Input Statistics:
  0 packets, 0 bytes
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts
  0 runts, 0 giants, 0 throttles
  0 CRC, 0 overrun, 0 discarded
Output Statistics:
  0 packets, 0 bytes, 0 underruns
  0 64-byte pkts, 0 over 64-byte pkts, 0 over 127-byte pkts
  0 over 255-byte pkts, 0 over 511-byte pkts, 0 over 1023-byte pkts
  0 Multicasts, 0 Broadcasts, 0 Unicasts
  0 throttles, 0 discarded, 0 collisions
Rate info (interval 299 seconds):
  Input 00.00 Mbits/sec,    0 packets/sec, 0.00% of line-rate
  Output 00.00 Mbits/sec,  0 packets/sec, 0.00% of line-rate
Time since last interface status change: 00:01:23

```

```

Dell#show running-config uplink-state-group
!
uplink-state-group 1
no enable
downstream TenGigabitEthernet 0/0
upstream TenGigabitEthernet 0/1
Dell#

```

```

Dell(conf-uplink-state-group-16)# show configuration
!
uplink-state-group 16
no enable
description test
downstream disable links all
downstream TengigabitEthernet 0/40
upstream TengigabitEthernet 0/41
upstream Port-channel 8

```

Sample Configuration: Uplink Failure Detection

The following example shows a sample configuration of UFD on a switch/router in which you configure as follows.

- Configure uplink-state group 3.
- Add downstream links GigabitEthernet 0/1, 0/2, 0/5, 0/9, 0/11, and 0/12.
- Configure two downstream links to be disabled if an upstream link fails.
- Add upstream links GigabitEthernet 0/3 and 0/4.
- Add a text description for the group.

- Verify the configuration with various show commands.

Example of Configuring UFD (S50)

```
Dell(conf)#uplink-state-group 3
Dell(conf-uplink-state-group-3)#

00:23:52: %STKUNIT0-M:CP %IFMGR-5-ASTATE_UP: Changed uplink state group Admin
state to up:
Group 3

Dell(conf-uplink-state-group-3)#downstream tengigabitethernet 0/1-2,5,9,11-12
Dell(conf-uplink-state-group-3)#downstream disable links 2
Dell(conf-uplink-state-group-3)#upstream tengigabitethernet 0/3-4
Dell(conf-uplink-state-group-3)#description Testing UFD feature
Dell(conf-uplink-state-group-3)#show config
!
uplink-state-group 3
  description Testing UFD feature
  downstream disable links 2
  downstream TenGigabitEthernet 0/1-2,5,9,11-12
  upstream TenGigabitEthernet 0/3-4

Dell#show running-config uplink-state-group
!
uplink-state-group 3
  description Testing UFD feature
  downstream disable links 2
  downstream TenGigabitEthernet 0/1-2,5,9,11-12
  upstream TenGigabitEthernet 0/3-4

Dell#show uplink-state-group 3

Uplink State Group: 3 Status: Enabled, Up

Dell#show uplink-state-group detail

(Up): Interface up (Dwn): Interface down (Dis): Interface disabled

Uplink State Group      : 3 Status: Enabled, Up
Upstream Interfaces     : Te 0/3(Up) Te 0/4(Up)
Downstream Interfaces   : Te 0/1(Up) Te 0/2(Up) Te 0/5(Up) Te 0/9(Up) Te 0/11(Up)
                        : Te 0/12(Up)

< After a single uplink port fails >

Dell#show uplink-state-group detail

(Up): Interface up (Dwn): Interface down (Dis): Interface disabled

Uplink State Group      : 3 Status: Enabled, Up
Upstream Interfaces     : Te 0/3(Dwn) Te 0/4(Up)
Downstream Interfaces   : Te 0/1(Dis) Te 0/2(Dis) Te 0/5(Up) Te 0/9(Up) Te
0/11(Up)
                        : Te 0/12(Up)
```

Upgrade Procedures

To find the upgrade procedures, go to the *Dell Networking OS Release Notes* for your system type to see all the requirements needed to upgrade to the desired Dell Networking OS version. To upgrade your system type, follow the procedures in the *Dell Networking OS Release Notes*.

Get Help with Upgrades

Direct any questions or concerns about the Dell Networking OS upgrade procedures to the Dell Technical Support Center. You can reach Technical Support:

- On the web: <http://www.dell.com/support/my-support/>
- By email: Dell-Force10_Technical_Support@Dell.com
- By phone: US and Canada: 866.965.5800, International: 408.965.5800.

Virtual LANs (VLANs)

Virtual LANs (VLANs) are supported on the MXL switch platform.

VLANs are a logical broadcast domain or logical grouping of interfaces in a local area network (LAN) in which all data received is kept locally and broadcast to all members of the group. When in Layer 2 mode, VLANs move traffic at wire speed and can span multiple devices. The Dell Networking operating system (OS) supports up to 4093 port-based VLANs and one default VLAN, as specified in IEEE 802.1Q.

VLANs benefits include:

- Improved security because you can isolate groups of users into different VLANs
- Ability to create one VLAN across multiple devices

For more information about VLANs, refer to the *IEEE Standard 802.1Q Virtual Bridged Local Area Networks*. In this guide, also refer to:

- [Bulk Configuration](#) in the [Interfaces](#) chapter.
- [VLAN Stacking](#) in the [Service Provider Bridging](#) chapter.

For a complete listing of all commands related to the Dell Networking OS VLANs, refer to these *Dell Networking OS Command Reference Guide* chapters:

- [Interfaces](#)
- [802.1X](#)
- [GARP VLAN Registration Protocol \(GVRP\)](#)
- [Service Provider Bridging](#)
- [Per-VLAN Spanning Tree Plus \(PVST+\)](#)

The following table lists the defaults for VLANs.

Feature	Default
Spanning Tree group ID	All VLANs are part of Spanning Tree group 0.
Mode	Layer 2 (no IP address is assigned).
Default VLAN ID	VLAN 1


Default VLAN

When you configure interfaces for Layer 2 mode, they are automatically placed in the Default VLAN as untagged interfaces. Only untagged interfaces can belong to the Default VLAN.

The following example displays the outcome of placing an interface in Layer 2 mode. To configure an interface for Layer 2 mode, use the `switchport` command. As shown in bold, the `switchport`

command places the interface in Layer 2 mode and the `show vlan` command in EXEC privilege mode indicates that the interface is now part of the Default VLAN (VLAN 1).

By default, VLAN 1 is the Default VLAN. To change that designation, use the `default vlan-id` command in CONFIGURATION mode. You cannot delete the Default VLAN.

 **NOTE:** You cannot assign an IP address to the Default VLAN. To assign an IP address to a VLAN that is currently the Default VLAN, create another VLAN and assign it to be the Default VLAN. For more information about assigning IP addresses, refer to [Assigning an IP Address to a VLAN](#).

- Untagged interfaces must be part of a VLAN. To remove an untagged interface from the Default VLAN, create another VLAN and place the interface into that VLAN. Alternatively, use the `no switchport` command, and Dell Networking OS removes the interface from the Default VLAN.
- A tagged interface requires an additional step to remove it from Layer 2 mode. Because tagged interfaces can belong to multiple VLANs, remove the tagged interface from all VLANs using the `no tagged interface` command. Only after the interface is untagged and a member of the Default VLAN can you use the `no switchport` command to remove the interface from Layer 2 mode. For more information, refer to [VLANs and Port Tagging](#).

Example of Configuring an Interface for Layer 2 Belonging to the Default VLAN

```
Dell(conf)#int tengig 3/2
Dell(conf-if)#no shut
Dell(conf-if)#switchport
Dell(conf-if)#show config
!
interface Tengigabitethernet 3/2
  no ip address
  switchport
  no shutdown
Dell(conf-if)#end
Dell#show vlan
```

Codes: * - Default VLAN, G - GVRP VLANs

```
NUM Status  Q Ports
* 1 Active   U Tengig 3/2
  2 Active   T Po1(So 0/0-1)
                    T Tengig 3/0
Dell#
```

Port-Based VLANs

Port-based VLANs are a broadcast domain defined by different ports or interfaces. In the Dell Networking OS, a port-based VLAN can contain interfaces from different line cards within the chassis. Dell Networking OS supports 4094 port-based VLANs.

Port-based VLANs offer increased security for traffic, conserve bandwidth, and allow switch segmentation. Interfaces in different VLANs do not communicate with each other, adding some security to the traffic on those interfaces. Different VLANs can communicate between each other by means of IP routing. Because traffic is only broadcast or flooded to the interfaces within a VLAN, the VLAN conserves bandwidth. Finally, you can have multiple VLANs configured on one switch, thus segmenting the device.

Interfaces within a port-based VLAN must be in Layer 2 mode and can be tagged or untagged in the VLAN ID.

VLANs and Port Tagging

To add an interface to a VLAN, the interface must be in Layer 2 mode. After you place an interface in Layer 2 mode, the interface is automatically placed in the Default VLAN.

The Dell Networking OS supports IEEE 802.1Q tagging at the interface level to filter traffic. When you enable tagging, a tag header is added to the frame after the destination and source MAC addresses. That information is preserved as the frame moves through the network. The following example shows the structure of a frame with a tag header. The VLAN ID is inserted in the tag header.



Figure 130. Tagged Frame Format

The tag header contains some key information that the Dell Networking OS uses:

- The VLAN protocol identifier identifies the frame as tagged according to the IEEE 802.1Q specifications (2 bytes).
- Tag control information (TCI) includes the VLAN ID (2 bytes total). The VLAN ID can have 4,096 values, but two are reserved.

NOTE: The insertion of the tag header into the Ethernet frame increases the size of the frame to more than the 1,518 bytes as specified in the IEEE 802.3 standard. Some devices that are not compliant with IEEE 802.3 may not support the larger frame size.

Information contained in the tag header allows the system to prioritize traffic and to forward information to ports associated with a specific VLAN ID. Tagged interfaces can belong to multiple VLANs, while untagged interfaces can belong only to one VLAN.

Configuration Task List

This section contains the following VLAN configuration tasks.


- [Creating a Port-Based VLAN](#) (mandatory)
- [Assigning Interfaces to a VLAN](#) (optional)
- [Assigning an IP Address to a VLAN](#) (optional)
- [Enabling Null VLAN as the Default VLAN](#) (optional)

Creating a Port-Based VLAN

To configure a port-based VLAN, create the VLAN and then add physical interfaces or port channel (LAG) interfaces to the VLAN.

NOTE: The Default VLAN (VLAN 1) is part of the system startup configuration and does not require configuration.

A VLAN is active only if the VLAN contains interfaces and those interfaces are operationally up. As shown in the following example, VLAN 1 is inactive because it does not contain any interfaces. The other VLANs contain enabled interfaces and are active.

 **NOTE:** In a VLAN, the `shutdown` command stops Layer 3 (routed) traffic only. Layer 2 traffic continues to pass through the VLAN. If the VLAN is not a routed VLAN (that is, configured with an IP address), the `shutdown` command has no affect on VLAN traffic.

When you delete a VLAN (using the `no interface vlan vlan-id` command), any interfaces assigned to that VLAN are assigned to the Default VLAN as untagged interfaces.

To create a port-based VLAN, use the following command.

- Configure a port-based VLAN (if the VLAN-ID is different from the Default VLAN ID) and enter INTERFACE VLAN mode.
CONFIGURATION mode

```
interface vlan vlan-id
```

To activate the VLAN, after you create a VLAN, assign interfaces in Layer 2 mode to the VLAN.

Example of Verifying a Port-Based VLAN

To view the configured VLANs, use the `show vlan` command in EXEC Privilege mode.

```
Dell#show vlan

Codes: * - Default VLAN, G - GVRP VLANs, R - Remote Port Mirroring VLANs, P -
Primary, C -
Community, I - Isolated
Q: U - Untagged, T - Tagged
   x - Dot1x untagged, X - Dot1x tagged
   G - GVRP tagged, M - Vlan-stack, H - VSN tagged
   i - Internal untagged, I - Internal tagged, v - VLT untagged, V - VLT tagged

NUM Status Description Q Ports
 1 Inactive a
* 20 Active           U Po32()
                       U Te 0/3,5,13,53-56
1002 Active          T Te 0/3,13,55-56
Dell#
```

Assigning Interfaces to a VLAN

You can only assign interfaces in Layer 2 mode to a VLAN using the tagged and untagged commands. To place an interface in Layer 2 mode, use the `switchport` command.

You can further designate these Layer 2 interfaces as tagged or untagged. For more information, refer to the [Interfaces](#) chapter and [Configuring Layer 2 \(Data Link\) Mode](#). When you place an interface in Layer 2 mode by the `switchport` command, the interface is automatically designated untagged and placed in the Default VLAN.

To view which interfaces are tagged or untagged and to which VLAN they belong, use the `show vlan` command. The following example shows that six VLANs are configured, and two interfaces are assigned to VLAN 2. The Q column in the `show vlan` command example notes whether the interface is tagged (T) or untagged (U). For more information about this command, refer to the Layer 2 chapter of the *Dell Networking OS Command Reference Guide*.

To tag frames leaving an interface in Layer 2 mode, assign that interface to a port-based VLAN to tag it with that VLAN ID. To tag interfaces, use the following commands.

1. Access INTERFACE VLAN mode of the VLAN to which you want to assign the interface.

CONFIGURATION mode

```
interface vlan vlan-id
```

2. Enable an interface to include the IEEE 802.1Q tag header.

INTERFACE mode

```
tagged interface
```

Add an Interface to Another VLAN

To view just the interfaces that are in Layer 2 mode, use the `show interfaces switchport` command in EXEC Privilege mode or EXEC mode.

The following example shows the steps to add a tagged interface (in this case, port channel 1) to VLAN 4. To view the interface's status. Interface (po 1) is tagged and in VLAN 2 and 3, use the `show vlan` command. In a port-based VLAN, use the `tagged` command to add the interface to another VLAN. The `show vlan` command output displays the interface's (po 1) changed status.

Except for hybrid ports, only a tagged interface can be a member of multiple VLANs. You can assign hybrid ports to two VLANs if the port is untagged in one VLAN and tagged in all others.

```
Dell#show vlan
```

Codes: * - Default VLAN, G - GVRP VLANs

```
NUM Status Q Ports
* 1 Inactive
  2 Active   T Pol(So 0/0-1)
              T Tengig 3/0
  3 Active   T Pol(So 0/0-1)
              T Tengig 3/1
```

```
Dell#config
Dell(conf)#int vlan 4
Dell(conf-if-vlan)#tagged po 1
Dell(conf-if-vlan)#show conf
!
interface Vlan 4
  no ip address
  tagged Port-channel 1
Dell(conf-if-vlan)#end
Dell#show vlan
```

Codes: * - Default VLAN, G - GVRP VLANs

```
NUM Status Q Ports
* 1 Inactive
  2 Active   T Pol(So 0/0-1)
              T Tengig 3/0
  3 Active   T Pol(So 0/0-1)
              T Tengig 3/1
  4 Active T Pol(So 0/0-1)
Dell#
```

When you remove a tagged interface from a VLAN (using the `no tagged interface` command), it remains tagged only if it is a tagged interface in another VLAN. If the tagged interface is removed from the only VLAN to which it belongs, the interface is placed in the Default VLAN as an untagged interface.

Moving Untagged Interfaces

To move untagged interfaces from the Default VLAN to another VLAN, use the following commands.

1. Access INTERFACE VLAN mode of the VLAN to which you want to assign the interface.
CONFIGURATION mode

```
interface vlan vlan-id
```

2. Configure an interface as untagged.
INTERFACE mode

```
untagged interface
```

This command is available only in VLAN interfaces.

Move an Untagged Interface to Another VLAN

The `no untagged interface` command removes the untagged interface from a port-based VLAN and places the interface in the Default VLAN. You cannot use the `no untagged interface` command in the Default VLAN. The following example shows the steps and commands to move an untagged interface from the Default VLAN to another VLAN.

To determine interface status, use the `show vlan` command. Interface (gi 3/2) is untagged and in the Default VLAN (vlan 1). In a port-based VLAN (vlan 4), use the `untagged` command to add the interface to that VLAN. The `show vlan` command output displays the interface's changed status (gi 3/2). Because the Default VLAN no longer contains any interfaces, it is listed as inactive.

```
Dell#show vlan

Codes: * - Default VLAN, G - GVRP VLANs
  NUM Status Q Ports
* 1 Active   U Tengig 3/2
  2 Active   T Po1(So 0/0-1)
                   T Tengig 3/0
  3 Active   T Po1(So 0/0-1)
                   T Tengig 3/1
  4 Inactive
Dell#conf
Dell(conf)#int vlan 4
Dell(conf-if-vlan)#untagged tengig 3/2
Dell(conf-if-vlan)#show config
!
interface Vlan 4
  no ip address
  untagged Tengigabitethernet 3/2
Dell(conf-if-vlan)#end
Dell#show vlan

Codes: * - Default VLAN, G - GVRP VLANs


  NUM Status Q Ports
* 1 Inactive
  2 Active   T Po1(So 0/0-1)
                   T Tengig 3/0
  3 Active   T Po1(So 0/0-1)
                   T Tengig 3/1
  4 Active   U Tengig 3/2
Dell#
```


The only way to remove an interface from the Default VLAN is to place the interface in Default mode by using the `no switchport` command in INTERFACE mode.

Assigning an IP Address to a VLAN

VLANs are a Layer 2 feature. For two physical interfaces on different VLANs to communicate, you must assign an IP address to the VLANs to route traffic between the two interfaces.

The `shutdown` command in INTERFACE mode does not affect Layer 2 traffic on the interface; the `shutdown` command only prevents Layer 3 traffic from traversing over the interface.

 **NOTE:** You cannot assign an IP address to the Default VLAN (VLAN 1). To assign another VLAN ID to the Default VLAN, use the `default vlan-id vlan-id` command.

In the Dell Networking OS, you can place VLANs and other logical interfaces in Layer 3 mode to receive and send routed traffic. For more information, refer to [Bulk Configuration](#).

To assign an IP address, use the following command.

- Configure an IP address and mask on the interface.
INTERFACE mode

```
ip address ip-address mask [secondary]
```

- `ip-address mask` — Enter an address in dotted-decimal format (A.B.C.D) and the mask must be in slash format (/24).
- `secondary` — This is the interface's backup IP address. You can configure up to eight secondary IP addresses.


Configuring Native VLANs

Traditionally, ports can be either untagged for membership to one VLAN or tagged for membership to multiple VLANs.

You must connect an untagged port to a VLAN-unaware station (one that does not understand VLAN tags), and you must connect a tagged port to a VLAN-aware station (one that generates and understands VLAN tags).

Native VLAN support breaks this barrier so that you can connect a port to both VLAN-aware and VLAN-unaware stations. Such ports are referred to as hybrid ports. Physical and port-channel interfaces may be hybrid ports.

Native VLAN is useful in deployments where a Layer 2 port can receive both tagged and untagged traffic on the same physical port. The classic example is connecting a voice-over-IP (VOIP) phone and a PC to the same port of the switch. The VOIP phone is configured to generate tagged packets (with VLAN = VOICE VLAN) and the attached PC generates untagged packets.

 **NOTE:** You cannot configure an existing switchport or port channel interface for Native VLAN. Interfaces must have no other Layer 2 or Layer 3 configurations when using the `portmode hybrid` command or a message similar to this displays: `% Error: Port is in Layer-2 mode Gi 5/6.`

To configure a port so that it can be a member of an untagged and tagged VLANs, use the following commands.

1. Remove any Layer 2 or Layer 3 configurations from the interface.

- INTERFACE mode
2. Configure the interface for Hybrid mode.
INTERFACE mode

`portmode hybrid`
 3. Configure the interface for Switchport mode.
INTERFACE mode

`switchport`
 4. Add the interface to a tagged or untagged VLAN.
VLAN INTERFACE mode

`[tagged | untagged]`

Enabling Null VLAN as the Default VLAN

In a Carrier Ethernet for Metro Service environment, service providers who perform frequent reconfigurations for customers with changing requirements occasionally enable multiple interfaces, each connected to a different customer, before the interfaces are fully configured.

This presents a vulnerability because both interfaces are initially placed in the native VLAN, VLAN 1, and for that period customers are able to access each other's networks. The Dell Networking OS has a Null VLAN to eliminate this vulnerability. When you enable the Null VLAN, all ports are placed into it by default, so even if you activate the physical ports of multiple customers, no traffic is allowed to traverse the links until each port is placed in another VLAN.

To enable Null VLAN, use the following command.

- Disable the default VLAN, so that all ports belong to the Null VLAN until configured as a member of another VLAN.

CONFIGURATION mode

```
default-vlan disable
```

Default: the default VLAN is enabled (no `default-vlan disable`).

Virtual Link Trunking (VLT)

Virtual link trunking (VLT) is supported on the MXL switch platform.

Overview

VLT allows physical links between two chassis to appear as a single virtual link to the network core.

VLT reduces the role of spanning tree protocols (STPs) by allowing link aggregation group (LAG) terminations on two separate distribution or core switches, and by supporting a loop-free topology. (To prevent the initial loop that may occur prior to VLT being established, use a spanning tree protocol. After VLT is established, you may use rapid spanning tree protocol (RSTP) to prevent loops from forming with new links that are incorrectly connected and outside the VLT domain.)

VLT peer devices have independent management planes. A chassis interconnect trunk between the VLT chassis maintains synchronization of L2/L3 control planes across the two VLT peers. The chassis interconnect trunk uses 10GE or 40GE user ports on the chassis.

VLT provides Layer 2 multipathing, creating redundancy through increased bandwidth, enabling multiple parallel paths between nodes and load-balancing traffic where alternative paths exist.

A separate backup link maintains heartbeat messages across an out-of-band management network. The backup link ensures that node failure conditions are correctly detected and are not confused with failures of the chassis interconnect trunk. VLT ensures that local traffic on a chassis does not traverse the chassis interconnect trunk and takes the shortest path to the destination via directly attached links.

Virtual link trunking offers the following benefits:

- Allows a single device to use a LAG across two upstream devices.
- Eliminates STP-blocked ports.
- Provides a loop-free topology.
- Uses all available uplink bandwidth.
- Provides fast convergence if either the link or a device fails.
- Optimized forwarding with virtual router redundancy protocol (VRRP).
- Provides link-level resiliency.
- Assures high availability.

As shown in the following example, VLT presents a single logical Layer 2 domain from the perspective of attached devices that have a virtual link trunk terminating on separate chassis in the VLT domain. However, the two VLT chassis are independent Layer2/Layer3 (L2/L3) switches for devices in the upstream network. L2/L3 control plane protocols and system management features function normally in VLT mode. Features such as VRRP and internet group management protocol (IGMP) snooping require state information coordinating between the two VLT chassis. IGMP and VLT configurations must be identical on both sides of the trunk to ensure the same behavior on both sides.

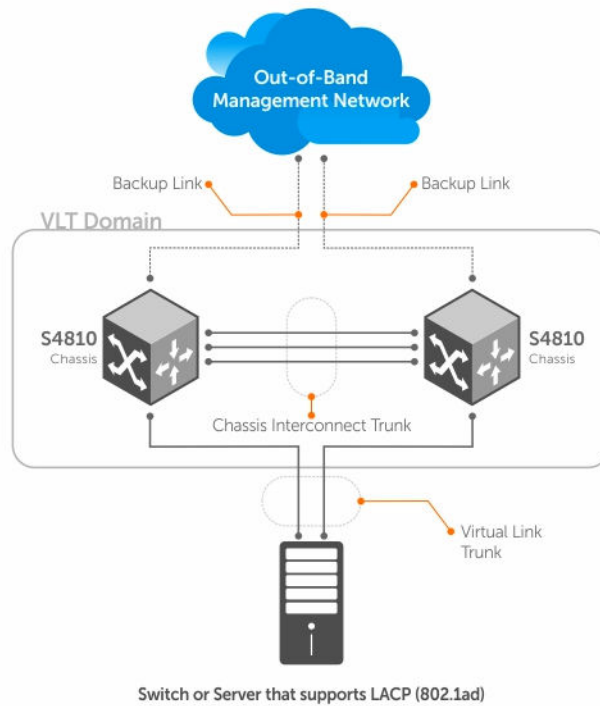


Figure 131. Virtual Link Trunking

Multi-domain VLT

A multi-domain VLT (mVLT) configuration creates a port channel between two VLT domains by allowing two different VLT domains, using different VLT Domain ID numbers, connected by a standard LACP LAG to form a loop-free Layer 2 topology in the aggregation layer.

This configuration supports a maximum of four (4) nodes per mVLT domain, increasing the number of available ports and allowing for dual redundancy of the VLT.

Additionally, a VLT domain that is a member of one mVLT can be used in another mVLT configuration with a different VLT domain. Routing protocols such as OSPF are not compatible with mVLT; however, VLT domains can be used for routing. A separate Layer 3 router is not required for inter-VLAN communication.

The following figure shows how the core/aggregation port density in the Layer 2 topology is increased using mVLT. For inter-VLAN routing and other Layer 3 routing, a separate Layer 3 router is required.

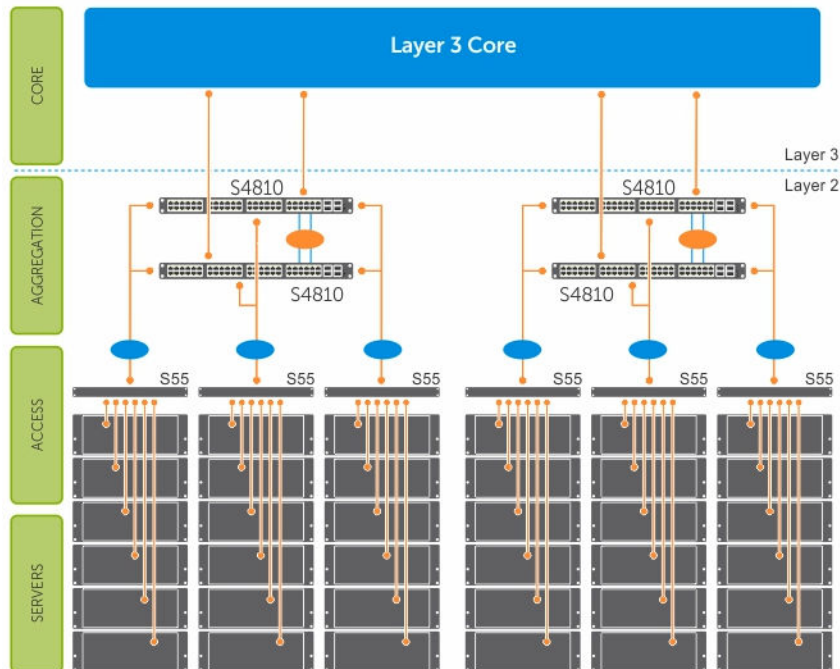


Figure 132. Multi-Domain VLT Example

VLT Terminology

The following are key VLT terms.

- **Virtual link trunk (VLT)** – The combined port channel between an attached device and the VLT peer switches.
- **VLT backup link** – The backup link monitors the vitality of VLT peer switches. The backup link sends configurable, periodic keep alive messages between the VLT peer switches.
- **VLT interconnect (VLTi)** – The link used to synchronize states between the VLT peer switches. Both ends must be on 10G or 40G interfaces.
- **VLT domain** – This domain includes both the VLT peer devices, VLT interconnect, and all of the port channels in the VLT connected to the attached devices. It is also associated to the configuration mode that you must use to assign VLT global parameters.
- **VLT peer device** – One of a pair of devices that are connected with the special port channel known as the VLT interconnect (VLTi).

Configure Virtual Link Trunking

VLT requires that you enable the feature and then configure the same VLT domain, backup link, and VLT interconnect on both peer switches.

Important Points to Remember

- VLT port channel interfaces must be switch ports.
- If you include RSTP on the system, configure it before VLT. Refer to [RSTP Configuration](#).
- Ensure that the spanning tree root bridge is at the Aggregation layer. If you enable RSTP on the VLT device, refer to [RSTP and VLT](#) for guidelines to avoid traffic loss.
- If you reboot both VLT peers in BMP mode and the VLT LAGs are static, the DHCP server reply to the DHCP discover offer may not be forwarded by the ToR to the correct node. To avoid this scenario, configure the VLT LAGs to the ToR and the ToR port channel to the VLT peers with LACP. If supported by the ToR, enable the `lacp-ungroup` feature on the ToR using the `lacp ungroup member-independent port-channel` command.
- If the `lacp-ungroup` feature is not supported on the ToR, reboot the VLT peers one at a time. After rebooting, verify that VLTi (ICL) is active before attempting DHCP connectivity.
- When you enable IGMP snooping on the VLT peers, ensure the value of the `delay-restore` command is not less than the query interval.
- When you enable Layer 3 routing protocols on VLT peers, make sure the `delay-restore` timer is set to a value that allows sufficient time for all routes to establish adjacency and exchange all the L3 routes between the VLT peers before you enable the VLT ports.
- Only use the `lacp ungroup member-independent` command if the system connects to nodes using bare metal provisioning (BMP) to upgrade or boot from the network.
- Ensure that you configure all port channels where LACP ungroup is applicable as hybrid ports and as untagged members of a VLAN. BMP uses untagged dynamic host configuration protocol (DHCP) packets to communicate with the DHCP server.
- If the DHCP server is located on the ToR and the VLTi (ICL) is down due to a failed link when a VLT node is rebooted in BMP mode, it is not able to reach the DHCP server, resulting in BMP failure.
- If the source is connected to an orphan (non-spanned, non-VLT) port in a VLT peer, the receiver is connected to a VLT (spanned) port-channel, and the VLT port-channel link between the VLT peer connected to the source and TOR is down, traffic is duplicated due to route inconsistency between peers. To avoid this scenario, Dell Networking recommends configuring both the source and the receiver on a spanned VLT VLAN.
- In a topology in which two VLT peer nodes that are connected by a VLTi link and are connected to a ToR switch using a VLT LAG interface, if you configure an egress IP ACL and apply it on the VLT LAG of both peers using the `deny ip any any` command, the traffic is permitted on the VLT LAG instead of being denied. The correct behavior of dropping the traffic on the VLT LAG occurs when VLT is up on both the peer nodes. However, if VLT goes down on one of the peers, traffic traverses through VLTi and the other peer switches it to the VLT LAG. Although egress ACL is applied on the VLT nodes to deny all traffic, this egress ACL does not deny the traffic (switching traffic is not denied owing to the egress IP ACL). You cannot use egress ACLs to deny traffic properly in such a VLT scenario.
- To support Q-in-Q over VLT, ICL is implicitly made as `vlan-stack trunk` port and the TPID of the ICL is set as **8100**.
- Layer 2 Protocol Tunneling is not supported in VLT.

Configuration Notes

When you configure VLT, the following conditions apply.

- VLT domain
 - A VLT domain supports two chassis members, which appear as a single logical device to network access devices connected to VLT ports through a port channel.

- A VLT domain consists of the two core chassis, the interconnect trunk, backup link, and the LAG members connected to attached devices.
- Each VLT domain has a unique MAC address that you create or VLT creates automatically.
- ARP tables are synchronized between the VLT peer nodes.
- VLT peer switches operate as separate chassis with independent control and data planes for devices attached on non-VLT ports.
- One chassis in the VLT domain is assigned a primary role; the other chassis takes the secondary role. The primary and secondary roles are required for scenarios when connectivity between the chassis is lost. VLT assigns the primary chassis role according to the lowest MAC address. You can configure the primary role.
- In a VLT domain, the peer switches must run the same Dell Networking operating system (OS) software version.
- Separately configure each VLT peer switch with the same VLT domain ID and the VLT version. If the system detects mismatches between VLT peer switches in the VLT domain ID or VLT version, the VLT Interconnect (VLTi) does not activate. To find the reason for the VLTi being down, use the `show vlt statistics` command to verify that there are mismatch errors, then use the `show vlt brief` command on each VLT peer to view the VLT version on the peer switch. If the VLT version is more than one release different from the current version in use, the VLTi does not activate.
- The chassis members in a VLT domain support connection to orphan hosts and switches that are not connected to both switches in the VLT core.
- VLT interconnect (VLTi)
 - The VLT interconnect must consist of either 10G or 40G ports. A maximum of sixteen 10G or 40G ports are supported. A combination of 10G and 40G ports are not supported.
 - A VLT interconnect over 1G ports is *not* supported.
 - The port channel must be in Default mode (not Switchport mode) to have VLTi recognize it.
 - The system automatically includes the required VLANs in VLTi. You do not need to manually select VLANs.
 - VLT peer switches operate as separate chassis with independent control and data planes for devices attached to non-VLT ports.
 - Port-channel link aggregation (LAG) across the ports in the VLT interconnect is required; individual ports are not supported. Dell Networking strongly recommends configuring a static LAG for VLTi.
 - IGMP state information is synchronized between the VLT chassis over the VLT interconnect.
 - The traffic transmitted over VLT interconnect is prioritized, and allows you to configure the traffic class-to-queue assignment.
 - The VLT interconnect synchronizes L2 and L3 control-plane information across the two chassis.
 - The VLT interconnect is used for data traffic only when there is a link failure that requires using VLTi in order for data packets to reach their final destination.
 - Unknown, multicast, and broadcast traffic can be flooded across the VLT interconnect.
 - MAC addresses for VLANs configured across VLT peer chassis are synchronized over the VLT interconnect on an egress port such as a VLT LAG. MAC addresses are the same on both VLT peer nodes.
 - ARP entries configured across the VLTi are the same on both VLT peer nodes.
 - If you shut down the port channel used in the VLT interconnect on a peer switch in a VLT domain in which you did not configure a backup link, the switch's role displays in the `show vlt brief` command output as Primary instead of Standalone.
 - When you change the default VLAN ID on a VLT peer switch, the VLT interconnect may flap.
 - In a VLT domain, the following software features are supported on VLTi: link layer discovery protocol (LLDP), flow control, port monitoring, jumbo frames, and data center bridging (DCB).
 - When you enable the VLTi link, the link between the VLT peer switches is established if the following configured information is true on both peer switches:

- * the VLT system MAC address matches.
- * the VLT unit-id is not identical.



NOTE: If you configure the VLT system MAC address or VLT unit-id on only one of the VLT peer switches, the link between the VLT peer switches is not established. Each VLT peer switch must be correctly configured to establish the link between the peers.

- If the link between the VLT peer switches is established, changing the VLT system MAC address or the VLT unit-id causes the link between the VLT peer switches to become disabled. However, removing the VLT system MAC address or the VLT unit-id may disable the VLT ports if you happen to configure the unit ID or system MAC address on only one VLT peer at any time.
 - If the link between VLT peer switches is established, any change to the VLT system MAC address or unit-id fails if the changes made create a mismatch by causing the VLT unit-ID to be the same on both peers and/or the VLT system MAC address does not match on both peers.
 - If you replace a VLT peer node, preconfigure the switch with the VLT system MAC address, unit-id, and other VLT parameters before connecting it to the existing VLT peer switch using the VLTi connection.
 - If the size of the MTU for VLTi members is less than 1496 bytes, MAC addresses may not be synced. Dell Networking recommends retaining the default MTU allocation (1554 bytes) for VLTi members.
- VLT backup link
 - In the backup link between peer switches, heartbeat messages are exchanged between the two chassis for health checks. The default time interval between heartbeat messages over the backup link is 1 second. You can configure this interval. The range is from 1 to 5 seconds. DSCP marking on heartbeat messages is CS6.
 - In order that the chassis backup link does not share the same physical path as the interconnect trunk, Dell Networking recommends using the management ports on the chassis and traverse an out-of-band management network. The backup link can use user ports, but not the same ports the interconnect trunk uses.
 - The chassis backup link does not carry control plane information or data traffic. Its use is restricted to health checks only.
 - Virtual link trunks (VLTs) between access devices and VLT peer switches
 - To connect servers and access switches with VLT peer switches, you use a VLT port channel, as shown in [Overview](#). Up to 96 port-channels are supported; up to 16 member links are supported in each port channel between the VLT domain and an access device.
 - The discovery protocol running between VLT peers automatically generates the ID number of the port channel that connects an access device and a VLT switch. The discovery protocol uses LACP properties to identify connectivity to a common client device and automatically generates a VLT number for port channels on VLT peers that connects to the device. The discovery protocol requires that an attached device always runs LACP over the port-channel interface.
 - VLT provides a loop-free topology for port channels with endpoints on different chassis in the VLT domain.
 - VLT uses shortest path routing so that traffic destined to hosts via directly attached links on a chassis does not traverse the chassis-interconnect link.
 - VLT allows multiple active parallel paths from access switches to VLT chassis.
 - VLT supports port-channel links with LACP between access switches and VLT peer switches. Dell Networking recommends using static port channels on VLTi.
 - If VLTi connectivity with a peer is lost but the VLT backup connectivity indicates that the peer is still alive, the VLT ports on the Secondary peer are orphaned and are shut down.
 - * In one possible topology, a switch uses the BMP feature to receive its IP address, configuration files, and boot image from a DHCP server that connects to the switch through the VLT domain. In the port-channel used by the switch to connect to the VLT domain, configure the port interfaces on each VLT peer as hybrid ports before adding them to the port channel (refer to

[Connecting a VLT Domain to an Attached Access Device \(Switch or Server\)](#)). To configure a port in Hybrid mode so that it can carry untagged, single-tagged, and double-tagged traffic, use the `portmode hybrid` command in Interface Configuration mode as described in [Configuring Native VLANs](#).

- * For example, if the DHCP server is on the ToR and VLTi (ICL) is down (due to either an unavailable peer or a link failure), whether you configured the VLT LAG as static or LACP, when a single VLT peer is rebooted in BMP mode, it cannot reach the DHCP server, resulting in BMP failure.
- Software features supported on VLT port-channels
 - In a VLT domain, the following software features are supported on VLT port-channels: 802.1p, ingress and egress ACLs, BGP, DHCP relay, IS-IS, OSPF, active-active PIM-SM, PIM-SSM, VRRP, Layer 3 VLANs, LLDP, flow control, port monitoring, jumbo frames, IGMP snooping, sFlow, ingress and egress ACLs, and Layer 2 control protocols RSTP only).
 - 🔑 **NOTE:** PVST+ passthrough is supported in a VLT domain. PVST+ BPDUs does not result in an interface shutdown. PVST+ BPDUs for a nondefault VLAN is flooded out as any other L2 multicast packet. On a default VLAN, RTSP is part of the PVST+ topology in that specific VLAN (default VLAN).
 - For detailed information about how to use VRRP in a VLT domain, refer to the following *VLT and VRRP Interoperability* section.
 - For information about configuring IGMP Snooping in a VLT domain, refer to [VLT and IGMP Snooping](#).
 - All system management protocols are supported on VLT ports, including SNMP, RMON, AAA, ACL, DNS, FTP, SSH, Syslog, NTP, RADIUS, SCP, TACACS+, Telnet, and LLDP.
 - Enable Layer 3 VLAN connectivity VLT peers by configuring a VLAN network interface for the same VLAN on both switches.
 - IGMP snooping is supported over VLT ports. The multicast forwarding state is synchronized on both VLT peer switches. The IGMP snooping process on a VLT peer shares the learned group information with the other VLT peer over the chassis interconnect trunk.
- Software features supported on VLT physical ports
 - In a VLT domain, the following software features are supported on VLT physical ports: 802.1p, LLDP, flow control, IPv6 dynamic routing, port monitoring, and jumbo frames.
- Software features not supported with VLT
 - In a VLT domain, the following software features are supported on non-VLT ports: 802.1x, DHCP snooping, FRRP, ingress and egress QOS.
- VLT and VRRP interoperability
 - In a VLT domain, VRRP interoperates with virtual link trunks that carry traffic to and from access devices (refer to [Overview](#)). The VLT peers belong to the same VRRP group and are assigned master and backup roles. Each peer actively forwards L3 traffic, reducing the traffic flow over the VLT interconnect.
 - VRRP elects the router with the highest priority as the master in the VRRP group. To ensure VRRP operation in a VLT domain, configure VRRP group priority on each VLT peer so that a peer is either the master or backup for all VRRP groups configured on its interfaces. For more information, refer to [Setting VRRP Group \(Virtual Router\) Priority](#).
 - To verify that a VLT peer is consistently configured for either the master or backup role in all VRRP groups, use the `show vrrp` command on each peer.
 - Also configure the same L3 routing (static and dynamic) on each peer so that the L3 reachability and routing tables are identical on both VLT peers. Both the VRRP master and backup peers must be able to locally forward L3 traffic in the same way.
 - In a VLT domain, although both VLT peers actively participate in L3 forwarding as the VRRP master or backup router, the `show vrrp` command output displays one peer as master and the other peer as backup.

- Failure scenarios
 - On a link failover, when a VLT port channel fails, the traffic destined for that VLT port channel is redirected to the VLTi to avoid flooding.
 - When a VLT switch determines that a VLT port channel has failed (and that no other local port channels are available), the peer with the failed port channel notifies the remote peer that it no longer has an active port channel for a link. The remote peer then enables data forwarding across the interconnect trunk for packets that would otherwise have been forwarded over the failed port channel. This mechanism ensures reachability and provides loop management. If the VLT interconnect fails, the VLT software on the primary switch checks the status of the remote peer using the backup link. If the remote peer is up, the secondary switch disables all VLT ports on its device to prevent loops.
 - If all ports in the VLT interconnect fail, or if the messaging infrastructure fails to communicate across the interconnect trunk, the VLT management system uses the backup link interface to determine whether the failure is a link-level failure or whether the remote peer has failed entirely. If the remote peer is still alive (heartbeat messages are still being received), the VLT secondary switch disables its VLT port channels. If keepalive messages from the peer are not being received, the peer continues to forward traffic, assuming that it is the last device available in the network. In either case, after recovery of the peer link or reestablishment of message forwarding across the interconnect trunk, the two VLT peers resynchronize any MAC addresses learned while communication was interrupted and the VLT system continues normal data forwarding.
 - If the primary chassis fails, the secondary chassis takes on the operational role of the primary.
- The SNMP MIB reports VLT statistics.

RSTP and VLT

VLT provides loop-free redundant topologies and does not require RSTP.

RSTP can cause temporary port state blocking and may cause topology changes after link or node failures. Spanning tree topology changes are distributed to the entire layer 2 network, which can cause a network-wide flush of learned MAC and ARP addresses, requiring these addresses to be re-learned. However, enabling RSTP can detect potential loops caused by non-system issues such as cabling errors or incorrect configurations. To minimize possible topology changes after link or node failure, RSTP is useful for potential loop detection. Configure RSTP using the following specifications.

The following recommendations help you avoid these issues and the associated traffic loss caused by using RSTP when you enable VLT on both VLT peers:

- Configure any ports at the edge of the spanning tree's operating domain as edge ports, which are directly connected to end stations or server racks. Disable RSTP on ports connected directly to Layer 3-only routers not running STP or configure them as edge ports.
- Ensure that the primary VLT node is the root bridge and the secondary VLT peer node has the second-best bridge ID in the network. If the primary VLT peer node fails, the secondary VLT peer node becomes the root bridge, avoiding problems with spanning tree port state changes that occur when a VLT node fails or recovers.
- Even with this configuration, if the node has non-VLT ports using RSTP that you did not configure as edge ports and are connected to other Layer 2 switches, spanning tree topology changes are still detected after VLT node recovery. To avoid this scenario, ensure that you configure any non-VLT ports as edge ports or disable RSTP.

VLT Bandwidth Monitoring

When bandwidth usage of the VLTi (ICL) exceeds 80%, a syslog error message (shown in the following message) and an SNMP trap are generated.

```
%STKUNIT0-M:CP %VLTMGR-6-VLT-LAG-ICL: Overall Bandwidth utilization of VLT-ICL-LAG (port-channel 25) crosses threshold. Bandwidth usage (80 )
```

When the bandwidth usage drops below the 80% threshold, the system generates another syslog message (shown in the following message) and an SNMP trap.

```
%STKUNIT0-M:CP %VLTMGR-6-VLT-LAG-ICL: Overall Bandwidth utilization of VLT-ICL-LAG (port-channel 25) reaches below threshold. Bandwidth usage (74 )VLT show remote port channel status
```

VLT and IGMP Snooping

When configuring IGMP Snooping with VLT, ensure the configurations on both sides of the VLT trunk are identical to get the same behavior on both sides of the trunk.

When you configure IGMP snooping on a VLT node, the dynamically learned groups and multicast router ports are automatically learned on the VLT peer node.

VLT Port Delayed Restoration

With the Dell Networking OS version 8.3.12.0, when a VLT node boots up, if the VLT ports have been previously saved in the start-up configuration, they are not immediately enabled.

To ensure MAC and ARP entries from the VLT per node are downloaded to the newly enabled VLT node, the system allows time for the VLT ports on the new node to be enabled and begin receiving traffic.

The `delay-restore` feature waits for all saved configurations to be applied, then starts a configurable timer. After the timer expires, the VLT ports are enabled one-by-one in a controlled manner. The delay between bringing up each VLT port-channel is proportional to the number of physical members in the port-channel. The default is 90 seconds.

To change the duration of the configurable timer, use the `delay-restore` command.

If you enable IGMP snooping, IGMP queries are also sent out on the VLT ports at this time allowing any receivers to respond to the queries and update the multicast table on the new node.

This delay in bringing up the VLT ports also applies when the VLTi link recovers from a failure that caused the VLT ports on the secondary VLT peer node to be disabled.

PIM-Sparse Mode Support on VLT

The designated router functionality of the PIM Sparse-Mode multicast protocol is supported on VLT peer switches for multicast sources and receivers that are connected to VLT ports.

VLT peer switches can act as a last-hop router for IGMP receivers and as a first-hop router for multicast sources.

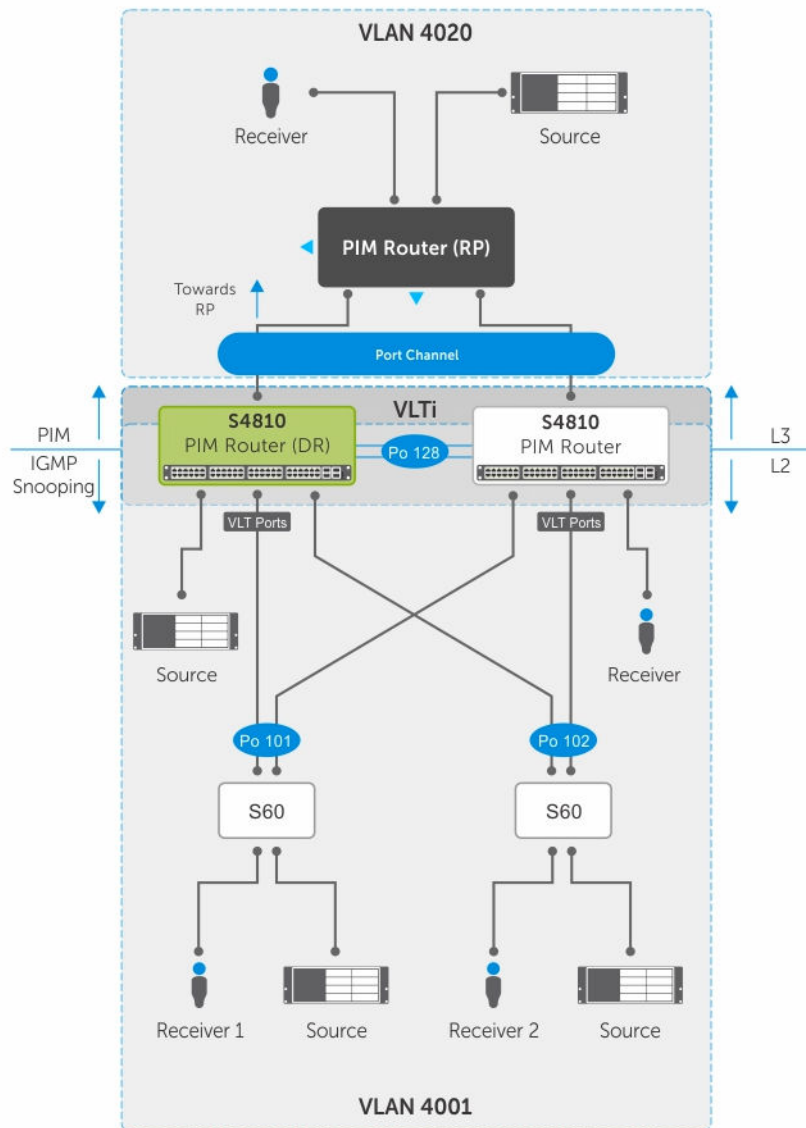


Figure 133. PIM-Sparse Mode Support on VLT

On each VLAN where the VLT peer nodes act as the first hop or last hop routers, one of the VLT peer nodes is elected as the PIM designated router. If you configured IGMP snooping along with PIM on the VLT VLANs, you must configure VLTi as the static multicast router port on both VLT peer switches. This ensures that for first hop routers, the packets from the source are redirected to the designated router

(DR) if they are incorrectly hashed. In addition to being first-hop or last-hop routers, the peer node can also act as an intermediate router.

The VLT peer nodes can also act as normal PIM routers on Layer 3 ports and on VLANs that do not have any VLT port members. In addition to being first-hop or last-hop routers, the peer node can also act as an intermediate router.

To route traffic to and from the multicast source and receiver that are connected to VLT ports, enable PIM-Sparse mode on the VLANs to which the VLT ports belong using the `ip pim sparse-mode` command. If IGMP Snooping is configured on these VLANs, the VLTi must be configured as a static multicast router port on both VLT peers.

To verify the PIM neighbors on the VLT VLAN and on the multicast port, use the `show ip pim neighbor`, `show ip igmp snooping mrouter`, and `show running config` commands.

You can configure virtual link trunking (VLT) peer nodes as rendezvous points (RPs) in a Protocol Independent Multicast (PIM) domain.

If the VLT node elected as the designated router fails, traffic loss occurs until another VLT node is elected the designated router.

VLT Multicast

VLT multicast provides multiple alternate paths for resiliency against link and node failures.

This feature supports inter-server multicast communication between top-of-rack (ToR) switches using an inter-VLAN Layer 3 routing protocol (for example, PIM, IS-IS, or OSPF). It also provides traffic resiliency during multicast routing convergence after failure without disrupting or altering multicast routing behavior.

Layer 2 protocols from the ToR to the server are intra-rack and inter-rack. No spanning tree is required, but interoperability with spanning trees at the aggregation layer is supported. Communication is active-active, with no blocked links. MAC tables are synchronized between VLT nodes for bridging and IGMP snooping can be enabled.

VLT multicast is also scalable, so additional racks can be implemented in an existing system to provide additional resiliency and resources to accommodate an increased need or to anticipate future growth. You can enable VLT unicast and multicast routing across multiple configurations using VLT links. Protocols such as BGP, OSPF, and PIM are compatible with VLT multicast.

Important Points to Remember

- You cannot configure a VLT node as a rendezvous point (RP), but any PIM-SM compatible VLT node can serve as a designated router (DR).
- You can only use one spanned VLAN from a PIM-enabled VLT node to an external neighboring PIM router.
- If you connect multiple spanned VLANs to a PIM neighbor, or if both spanned and non-spanned VLANs can access the PIM neighbor, ECMP can cause the PIM protocol running on each VLT peer node to choose a different VLAN or IP route to reach the PIM neighbor. This can result in issues with multicast route syncing between peers.
- Both VLT peers require symmetric Layer 2 and Layer 3 configurations on both VLT peers for any spanned VLAN.

- For optimal performance, configure the VLT VLAN routing metrics to prefer VLT VLAN interfaces over non-VLT VLAN interfaces.
- When using factory default settings on a new switch deployed as a VLT node, packet loss may occur due to the requirement that all ports must be open.
- You can enable ECMP on VLT nodes using VLT unicast; however, ECMP is not compatible on VLT nodes using VLT multicast. You must use a single VLAN.

Configuring VLT Multicast

To enable and configure VLT multicast, follow these steps.

1. Enable VLT on a switch, then configure a VLT domain and enter VLT-domain configuration mode.

CONFIGURATION mode

```
vlt domain domain-id
```

2. Enable peer-routing.

VLT DOMAIN mode

```
peer-routing
```

3. Configure the multicast peer-routing timeout.

VLT DOMAIN mode

```
multicast peer-routing-timeout value
```

value: Specify a value (in seconds) from 1 to 1200.

4. Configure a PIM-SM compatible VLT node as a designated router (DR). For more information, refer to [Configuring a Designated Router](#).
5. Configure a PIM-enabled external neighboring router as a rendezvous point (RP). For more information, refer to [Configuring a Static Rendezvous Point](#).
6. Configure the VLT VLAN routing metrics to prefer VLT VLAN interfaces over non-VLT VLAN interfaces. For more information, refer to [Classify Traffic](#).
7. Configure symmetrical Layer 2 and Layer 3 configurations on both VLT peers for any spanned VLAN.

VLT Unicast Routing

VLT unicast locally routes packets destined for the L3 endpoint of the VLT peer.

This method avoids sub-optimal routing. Peer-routing syncs the MAC addresses of both VLT peers and requires two local DA entries in TCAM. In case a VLT node is down, resiliency is provided by a timer that allows you to configure the amount of time needed for peer recovery.


VLT unicast also assists in conflict resolution by updating the routing table when a path for an existing multicast route has been changed. You can enable VLT unicast and multicast routing across multiple configurations using VLT links. ECMP can be enabled on VLT nodes using VLT unicast; however, ECMP is not compatible on VLT nodes using VLT multicast. You must use a single VLAN.


VLT unicast is supported on both IPV6 / IPv4 . To enable VLT unicast, both VLT peers must be in L3 mode. Static route and routing protocols such as RIP, OSPF, ISIS, and BGP are supported. However, point-to-point configuration is not supported. To enable VLT unicast, VLAN configuration must be symmetrical on both peers. The same VLAN cannot be configured as Layer 2 on one node and as Layer 3 on the other

node. Configuration mismatches are logged in the syslog and displayed in the output of the `show vlt inconsistency` command.

When you enable VLT unicast, VLAN wildcarding is enabled to support up to 4094 VLANs. If you enable VLT unicast, the following actions occur:

- L3 routing is enabled on any new IP address / IPV6 address configured for a VLAN interface that is up.
- L3 routing is enabled on any VLAN with an admin state of up.
- For PVLAN, if the IP address is configured for the primary VLAN, L3 routing is enabled.

 **NOTE:** If the CAM is full, do not enable peer-routing.

 **NOTE:** The peer routing and peer-routing-timeout is applicable for both IPv6/ IPv4.

Configuring VLT Unicast

To enable and configure VLT unicast, follow these steps.

1. Enable VLT on a switch, then configure a VLT domain and enter VLT-domain configuration mode.
CONFIGURATION mode

```
vlt domain domain-id
```

2. Enable peer-routing.
VLT DOMAIN mode

```
peer-routing
```

3. Configure the peer-routing timeout.
VLT DOMAIN mode

```
peer-routing-timeout value
```

value: Specify a value (in seconds) from 1 to 65535.

Non-VLT ARP Sync

In the Dell Networking OS version 9.2(0.0), ARP entries (including ND entries) learned on other ports are synced with the VLT peer to support station move scenarios.

Prior to Dell Networking OS version 9.2.(0.0), only ARP entries learned on VLT ports were synced between peers.

Additionally, ARP entries resulting from station movements from VLT to non-VLT ports or to different non-VLT ports are learned on the non-VLT port and synced with the peer node. The peer node is updated to use the new non-VLT port.

 **NOTE:** ARP entries learned on non-VLT, non-spanned VLANs are not synced with VLT peers.

RSTP Configuration

RSTP is supported in a VLT domain.

Before you configure VLT on peer switches, configure RSTP in the network. RSTP is required for initial loop prevention during the VLT startup phase. You may also use RSTP for loop prevention in the network outside of the VLT port channel. For information about how to configure RSTP, [Rapid Spanning Tree Protocol \(RSTP\)](#).

Run RSTP on both VLT peer switches. The primary VLT peer controls the RSTP states, such as forwarding and blocking, on both the primary and secondary peers. Dell Networking recommends configuring the primary VLT peer as the RSTP primary root device and configuring the secondary VLT peer as the RSTP secondary root device.

BPDUs use the MAC address of the primary VLT peer as the RSTP bridge ID in the designated bridge ID field. The primary VLT peer sends these BPDUs on VLT interfaces connected to access devices. The MAC address for a VLT domain is automatically selected on the peer switches when you create the domain.

Configure both ends of the VLT interconnect trunk with identical RSTP configurations. When you enable VLT, the `show spanning-tree rstp brief` command output displays VLT information.

Preventing Forwarding Loops in a VLT Domain

During the bootup of VLT peer switches, a forwarding loop may occur until the VLT configurations are applied on each switch and the primary/secondary roles are determined.

To prevent the interfaces in the VLT interconnect trunk and RSTP-enabled VLT ports from entering a Forwarding state and creating a traffic loop in a VLT domain, take the following steps.

1. Configure RSTP in the core network and on each peer switch as described in [Rapid Spanning Tree Protocol \(RSTP\)](#).

Disabling RSTP on one VLT peer may result in a VLT domain failure.

2. Enable RSTP on each peer switch.

```
PROTOCOL SPANNING TREE RSTP mode
```

```
no disable
```

3. Configure each peer switch with a unique bridge priority.

```
PROTOCOL SPANNING TREE RSTP mode
```

```
bridge-priority
```

Sample RSTP Configuration

The following is a sample of an RSTP configuration.

Using the example shown in the [Protocol Overview](#) section as a sample VLT topology, the primary VLT switch sends BPDUs to an access device (switch or server) with its own RSTP bridge ID. BPDUs generated by an RSTP-enabled access device are only processed by the primary VLT switch. The secondary VLT switch tunnels the BPDUs that it receives to the primary VLT switch over the VLT interconnect. Only the

primary VLT switch determines the RSTP roles and states on VLT ports and ensures that the VLT interconnect link is never blocked.

In the case of a primary VLT switch failure, the secondary switch starts sending BPDUs with its own bridge ID and inherits all the port states from the last synchronization with the primary switch. An access device never detects the change in primary/secondary roles and does not see it as a topology change.

The following examples show the RSTP configuration that you must perform on each peer switch to prevent forwarding loops.

Configure RSTP on VLT Peers to Prevent Forwarding Loops (VLT Peer 1)

```
Dell_VLTpeer1(conf)#protocol spanning-tree rstp
Dell_VLTpeer1(conf-rstp)#no disable
Dell_VLTpeer1(conf-rstp)#bridge-priority 4096
```


Configure RSTP on VLT Peers to Prevent Forwarding Loops (VLT Peer 2)

```
Dell_VLTpeer2(conf)#protocol spanning-tree rstp
Dell_VLTpeer2(conf-rstp)#no disable
Dell_VLTpeer2(conf-rstp)#bridge-priority 0
```

Configuring VLT

To configure virtual link trunking and create a VLT domain in which two MXL switches are physically connected and treated as a single port channel by access devices, you must configure the following settings on each VLT peer device.

Prerequisites: Before you begin, make sure that both VLT peer switches are running the same Dell Networking OS version and are configured for RSTP as described in [RSTP Configuration](#). For VRRP operation, ensure that you configure VRRP groups and L3 routing on each VLT peer as described in *VLT and VRRP interoperability* in the [Configuration Notes](#) section.

1. Configure the VLT interconnect for the VLT domain. The primary and secondary switch roles in the VLT domain are automatically assigned after you configure both sides of the VLTi.
 -  **NOTE:** If you use a third-party ToR unit, to avoid potential problems if you reboot the VLT peers, Dell recommends using static LAGs on the VLTi between VLT peers.
2. Enable VLT and create a VLT domain ID. VLT automatically selects a system MAC address.
3. Configure a backup link for the VLT domain.
4. (Optional) Manually reconfigure the default VLT settings, such as the MAC address and VLT primary/secondary roles.
5. Connect the peer switches in a VLT domain to an attached access device (switch or server).

Configuring a VLT Interconnect


To configure a VLT interconnect, follow these steps.

1. Configure the port channel for the VLT interconnect on a VLT switch and enter interface configuration mode.

CONFIGURATION mode

```
interface port-channel id-number
```

Enter the same port-channel number configured with the `peer-link port-channel` command as described in [Configuring VLT](#) and [Connecting a VLT Domain](#).

 **NOTE:** To be included in the VLTi, the port channel must be in Default mode (no `switchport` or VLAN assigned).

2. Remove an IP address from the interface.

INTERFACE PORT-CHANNEL mode

```
no ip address
```

3. Add one or more port interfaces to the port channel.

INTERFACE PORT-CHANNEL mode

```
channel-member interface
```

interface: specify one of the following interface types:

- 10-Gigabit Ethernet: Enter `tengigabitethernet slot/port`.
- 40-Gigabit Ethernet: Enter `fortygigabitethernet slot/port`.

4. Ensure that the port channel is active.

INTERFACE PORT-CHANNEL mode

```
no shutdown
```

5. Repeat Steps 1 to 4 on the VLT peer switch to configure the VLT interconnect.

Configuring a VLT Backup Link

To configure a VLT backup link, use the following command.

1. Specify the management interface to be used for the backup link through an out-of-band management network.

CONFIGURATION mode

```
interface managementethernet slot/ port
```

Enter the slot (0-1) and the port (0).

2. Configure an IPv4 address (A.B.C.D) or IPv6 address (X:X:X::X) and mask (/x) on the interface.

MANAGEMENT INTERFACE mode

```
{ip address ipv4-address/ mask | ipv6 address ipv6-address/ mask}
```

This is the IP address to be configured on the VLT peer with the `back-up destination` command.

3. Ensure that the interface is active.

MANAGEMENT INTERFACE mode

```
no shutdown
```

4. Repeat Steps 1 to 3 on the VLT peer switch.

To set an amount of time, in seconds, to delay the system from restoring the VLT port, use the `delay-restore` command at any time. For more information, refer to [VLT Port Delayed Restoration](#).

Configuring a VLT Port Delay Period

To configure a VLT port delay period, use the following commands.

1. Enter VLT-domain configuration mode for a specified VLT domain.

CONFIGURATION mode

```
vlt domain domain-id
```

The range of domain IDs from 1 to 1000.

2. Enter an amount of time, in seconds, to delay the restoration of the VLT ports after the system is rebooted.

CONFIGURATION mode

```
delay-restore delay-restore-time
```

The range is from 1 to 1200.

The default is **90 seconds**.

Reconfiguring the Default VLT Settings (Optional)

To reconfigure the default VLT settings, use the following commands.

1. Enter VLT-domain configuration mode for a specified VLT domain.

CONFIGURATION mode

```
vlt domain domain-id
```

The range of domain IDs is from 1 to 1000.

2. (Optional) After you configure the VLT domain on each peer switch on both sides of the interconnect trunk, by default, the system elects a primary and secondary VLT peer device.

VLT DOMAIN CONFIGURATION mode

```
primary-priority value
```

To reconfigure the primary role of VLT peer switches, use the `primary-priority` command. To configure the primary role on a VLT peer, enter a lower value than the priority value of the remote peer.

The priority values are from 1 to 65535. The default is **32768**.

3. (Optional) When you create a VLT domain on a switch, the system automatically creates a VLT-system MAC address used for internal system operations.

VLT DOMAIN CONFIGURATION mode

```
system-mac mac-address mac-address
```

To explicitly configure the default MAC address for the domain by entering a new MAC address, use the `system-mac` command. The format is `aaaa.bbbb.cccc`.

Also, reconfigure the same MAC address on the VLT peer switch.

Use this command to minimize the time required for the VLT system to synchronize the default MAC address of the VLT domain on both peer switches when one peer switch reboots.

4. (Optional) When you create a VLT domain on a switch, the system automatically assigns a unique unit ID (0 or 1) to each peer switch.

VLT DOMAIN CONFIGURATION mode

```
unit-id {0 | 1}
```

To explicitly configure the default values on each peer switch, use the `unit-id` command.

Configure a different unit ID (0 or 1) on each peer switch.

Unit IDs are used for internal system operations.

Use this command to minimize the time required for the VLT system to determine the unit ID assigned to each peer switch when one peer switch reboots.

Connecting a VLT Domain to an Attached Access Device (Switch or Server)

To connect a VLT domain to an attached access device, use the following commands.

On a VLT peer switch: To connect to an attached device, configure the same port channel ID number on each peer switch in the VLT domain.

1. Configure the same port channel to be used to connect to an attached device and enter interface configuration mode.

CONFIGURATION mode

```
interface port-channel id-number
```

2. Remove an IP address from the interface.

INTERFACE PORT-CHANNEL mode

```
no ip address
```

3. Place the interface in Layer 2 mode.

INTERFACE PORT-CHANNEL mode

```
switchport
```

4. Add one or more port interfaces to the port channel.

INTERFACE PORT-CHANNEL mode

```
channel-member interface
```

interface: specify one of the following interface types:

- 10-Gigabit Ethernet: enter `tengigabitethernet slot/port`.
- 40-Gigabit Ethernet: enter `fortygigabitethernet slot/port`.

5. Ensure that the port channel is active.

INTERFACE PORT-CHANNEL mode

```
no shutdown
```

6. Associate the port channel to the corresponding port channel in the VLT peer for the VLT connection to an attached device.

INTERFACE PORT-CHANNEL mode

```
vlt-peer-lag port-channel id-number
```

The valid port-channel ID numbers are from 1 to 128.

7. Repeat Steps 1 to 6 on the VLT peer switch to configure the same port channel as part of the VLT domain.
8. **On an attached switch or server:** To connect to the VLT domain and add port channels to it, configure a port channel.

To configure the VLAN where a VLT peer forwards received packets over the VLTi from an adjacent VLT peer that is down, use the `peer-down-vlan` parameter. When a VLT peer with BMP reboots, untagged DHCP discover packets are sent to the peer over the VLTi. Using this configuration ensures the DHCP discover packets are forwarded to the VLAN that has the DHCP server.

Configuring a VLT VLAN Peer-Down (Optional)

To configure a VLT VLAN peer-down, use the following commands.

1. Enter VLT-domain configuration mode for a specified VLT domain.

CONFIGURATION mode

```
vlt domain domain-id
```

The range of domain IDs is from 1 to 1000.

2. Enter the port-channel number that acts as the interconnect trunk.

VLT DOMAIN CONFIGURATION mode

```
peer-link port-channel id-number
```

The range is from 1 to 128.

3. Enter the VLAN ID number of the VLAN where the VLT forwards packets received on the VLTi from an adjacent peer that is down.

VLT DOMAIN CONFIGURATION mode

```
peer-down-vlan vlan interface number
```

The range is from 1 to 4094.

Configure Multi-domain VLT (mVLT) (Optional)

To configure a multi-domain VLT between two VLT domains on your network, use the following procedure.

For a sample configuration, refer to the [Multi-domain VLT](#) section.

1. Configure the port channel to be used for the VLT interconnect on a VLT switch and enter Interface Configuration mode.

CONFIGURATION mode

```
interface port-channel id-number
```

Enter the same port-channel number configured with the `peer-link port-channel` command.

2. Add one or more port interfaces to the port channel.

INTERFACE PORT-CHANNEL mode

```
channel-member interface
```

interface specifies one of the following interface types:

- 10-Gigabit Ethernet: Enter `tengigabitethernet slot/port`.
- 40-Gigabit Ethernet: Enter `fortygigabitethernet slot/port`.

3. Enter VLT-domain configuration mode for a specified VLT domain.

CONFIGURATION mode

```
vlt domain domain-id
```

The range of domain IDs is from 1 to 1000.

4. Enter the port-channel number that will act as the interconnect trunk.

VLT DOMAIN CONFIGURATION mode

```
peer-link port-channel id-number
```

The range is from 1 to 128.

5. Configure the IP address of the management interface on the remote VLT peer to be used as the endpoint of the VLT backup link for sending out-of-band hello messages.

VLT DOMAIN CONFIGURATION mode

```
back-up destination ip-address [interval seconds]
```

You can optionally specify the time interval used to send hello messages.

The range is from 1 to 5 seconds.

6. When you create a VLT domain on a switch, the system automatically creates a VLT-system MAC address used for internal system operations.

VLT DOMAIN CONFIGURATION mode

```
system-mac mac-address mac-address
```

Use the `system-mac` command to explicitly configure the default MAC address for the domain by entering a new MAC address in the format: `aaaa.bbbb.cccc`.

You must also reconfigure the same MAC address on the VLT peer switch.

Use this command to minimize the time required for the VLT system to synchronize the default MAC address of the VLT domain on both peer switches when one peer switch reboots.

7. When you create a VLT domain on a switch, the system automatically assigns a unique unit ID (0 or 1) to each peer switch. The unit IDs are used for internal system operations.

VLT DOMAIN CONFIGURATION mode

```
unit-id {0 | 1}
```

Use the `unit-id` command to explicitly configure the default values on each peer switch.

You must configure a different unit ID (0 or 1) on each peer switch.

Use this command to minimize the time required for the VLT system to determine the unit ID assigned to each peer switch when one peer switch reboots.

- 8. Configure multi-domain VLT.** Configure the port channel to be used for the VLT interconnect on a VLT switch and enter interface configuration mode.

CONFIGURATION mode

```
interface port-channel id-number
```

Enter the same port-channel number configured with the `peer-link port-channel` command.

- 9.** Place the interface in Layer 2 mode.

INTERFACE PORT-CHANNEL mode

```
switchport
```

- 10.** Associate the port channel to the corresponding port channel in the VLT peer for the VLT connection to an attached device.

INTERFACE PORT-CHANNEL mode

```
vlt-peer-lag port-channel id-number
```

Valid port-channel ID numbers are from 1 to 128.

- 11.** Ensure that the port channel is active.

INTERFACE PORT-CHANNEL mode

```
no shutdown
```

- 12. Add links to the mVLT port.** Configure a range of interfaces to bulk configure.

CONFIGURATION mode

```
interface range {port-channel id}
```

- 13.** Enable LACP on the LAN port.

INTERFACE mode

```
port-channel-protocol lacp
```

- 14.** Configure the LACP port channel mode.

INTERFACE mode

```
port-channel number mode [active]
```

- 15.** Ensure that the interface is active.

MANAGEMENT INTERFACE mode

```
no shutdown
```

- 16.** Repeat steps 1 through 15 for the VLT peer node in Domain 1.

- 17.** Repeat steps 1 through 15 for the first VLT node in Domain 2.

- 18.** Repeat steps 1 through 15 for the VLT peer node in Domain 2.

Verifying a VLT Configuration

To monitor the operation or verify the configuration of a VLT domain, use any of the following `show` commands on the primary and secondary VLT switches.

- Display information on backup link operation.
EXEC mode

`show vlt backup-link`
- Display general status information about VLT domains currently configured on the switch.
EXEC mode

`show vlt brief`
- Display detailed information about the VLT-domain configuration, including local and peer port-channel IDs, local VLT switch status, and number of active VLANs on each port channel.
EXEC mode

`show vlt detail`
- Display the VLT peer status, role of the local VLT switch, VLT system MAC address and system priority, and the MAC address and priority of the locally-attached VLT device.
EXEC mode

`show vlt role`
- Display the current configuration of all VLT domains or a specified group on the switch.
EXEC mode

`show running-config vlt`
- Display statistics on VLT operation.
EXEC mode

`show vlt statistics`
- Display the RSTP configuration on a VLT peer switch, including the status of port channels used in the VLT interconnect trunk and to connect to access devices.
EXEC mode

`show spanning-tree rstp`
- Display the current status of a port or port-channel interface used in the VLT domain.
EXEC mode

`show interfaces interface`
 - *interface*: specify one of the following interface types:
 - * Fast Ethernet: enter `fastethernet slot/port`.
 - * 10-Gigabit Ethernet: enter `tengigabitethernet slot/port`.
 - * 40-Gigabit Ethernet: enter `fortygigabitethernet slot/port`.
 - * Port channel: enter `port-channel {1-128}`.

Example of the show vlt backup-link Command

Example of the show vlt brief Command

Example of the show vlt detail Command

Example of the show vlt role Command

Example of the show running-config vlt Command

Example of the show vlt statistics Command

Example of the show spanning-tree rstp Command

```
Dell_VLTpeer1# show vlt backup-link
```

```
VLT Backup Link
-----
Destination:          10.11.200.18
Peer HeartBeat status: Up
HeartBeat Timer Interval: 1
HeartBeat Timeout:    3
UDP Port:             34998
HeartBeat Messages Sent: 1026
HeartBeat Messages Received: 1025
```

```
Dell_VLTpeer2# show vlt backup-link
```

```
VLT Backup Link
-----
Destination:          10.11.200.20
Peer HeartBeat status: Up
HeartBeat Timer Interval: 1
HeartBeat Timeout:    3
UDP Port:             34998
HeartBeat Messages Sent: 1030
HeartBeat Messages Received: 1014
```

```
Dell#show vlt br
VLT Domain Brief
```

```
-----
Domain ID                : 1
Role                     : Secondary
Role Priority             : 32768
ICL Link Status          : Up
HeartBeat Status         : Up
VLT Peer Status          : Up
Version                  : 6(3)
Local System MAC address : 00:01:e8:8a:e9:91
Remote System MAC address : 00:01:e8:8a:e9:76
Remote system version    : 6(3)
Delay-Restore timer      : 90 seconds

Delay-Restore Abort Threshold : 60 seconds
Peer-Routing                : Disabled
Peer-Routing-Timeout timer  : 0 seconds
Multicast peer-routing timeout : 150 seconds
Dell#
```

```
Dell_VLTpeer1# show vlt detail
```

```
Local LAG Id Peer LAG Id Local Status Peer Status Active VLANs
-----
100          100          UP          UP          10, 20, 30
```

```
127          2          UP          UP          20, 30
```

```
Dell_VLTpeer2# show vlt detail
```

Local LAG Id	Peer LAG Id	Local Status	Peer Status	Active VLANs
2	127	UP	UP	20, 30
100	100	UP	UP	10, 20, 30

```
Dell_VLTpeer1# show vlt role
```

```
VLT Role
```

```
-----  
VLT Role: Primary  
System MAC address: 00:01:e8:8a:df:bc  
System Role Priority: 32768  
Local System MAC address: 00:01:e8:8a:df:bc  
Local System Role Priority: 32768
```

```
Dell_VLTpeer2# show vlt role
```

```
VLT Role
```

```
-----  
VLT Role: Secondary  
System MAC address: 00:01:e8:8a:df:bc  
System Role Priority: 32768  
Local System MAC address: 00:01:e8:8a:df:e6  
Local System Role Priority: 32768
```

```
Dell_VLTpeer1# show running-config vlt
```

```
!  
vlt domain 30  
  peer-link port-channel 60  
  back-up destination 10.11.200.18
```

```
Dell_VLTpeer2# show running-config vlt
```

```
!  
vlt domain 30  
  peer-link port-channel 60  
  back-up destination 10.11.200.20
```

```
Dell_VLTpeer1# show vlt statistics
```

```
VLT Statistics
```

```
-----  
HeartBeat Messages Sent: 987  
HeartBeat Messages Received: 986  
ICL Hello's Sent: 148  
ICL Hello's Received: 98
```

```
Dell_VLTpeer2# show vlt statistics
```

```
VLT Statistics
```

```
-----  
HeartBeat Messages Sent: 994  
HeartBeat Messages Received: 978  
ICL Hello's Sent: 89  
ICL Hello's Received: 89
```

The bold section displays the RSTP state of port channels in the VLT domain. Port channel 100 is used in the VLT interconnect trunk (VLTi) to connect to VLT peer2. Port channels 110, 111, and 120 are used to connect to access switches or servers (vlt).

```
Dell_VLTpeer1# show spanning-tree rstp brief
```

```
Executing IEEE compatible Spanning Tree Protocol
Root ID Priority 0, Address 0001.e88a.dff8
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID Priority 4096, Address 0001.e88a.d656
Configured hello time 2, max age 20, forward delay 15
```

Interface Name	PortID	Prio	Cost	Sts	Cost	Designated Bridge ID	PortID
Po 1	128.2	128	200000	DIS	800	4096	0001.e88a.d656 128.2
Po 3	128.4	128	200000	DIS	800	4096	0001.e88a.d656 128.4
Po 4	128.5	128	200000	DIS	800	4096	0001.e88a.d656 128.5
Po 100	128.101	128	800	FWD (VLTi)	800	0	0001.e88a.dff8 128.101
Po 110	128.111	128	00	FWD (vlt)	800	4096	0001.e88a.d656 128.111
Po 111	128.112	128	200000	DIS (vlt)	800	4096	0001.e88a.d656 128.112
Po 120	128.121	128	2000	FWD (vlt)	800	4096	0001.e88a.d656 128.121

```
Dell_VLTpeer2# show spanning-tree rstp brief
```

```
Executing IEEE compatible Spanning Tree Protocol
Root ID Priority 0, Address 0001.e88a.dff8
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID Priority 0, Address 0001.e88a.dff8
We are the root
Configured hello time 2, max age 20, forward delay 15
```

Interface Name	PortID	Prio	Cost	Sts	Cost	Designated Bridge ID	PortID
Po 1	128.2	128	200000	DIS	0	0	0001.e88a.dff8 128.2
Po 3	128.4	128	200000	DIS	0	0	0001.e88a.dff8 128.4
Po 4	128.5	128	200000	DIS	0	0	0001.e88a.dff8 128.5
Po 100	128.101	128	800	FWD (VLTi)	0	0	0001.e88a.dff8 128.101
Po 110	128.111	128	00	FWD (vlt)	0	0	0001.e88a.dff8 128.111
Po 111	128.112	128	200000	DIS (vlt)	0	0	0001.e88a.dff8 128.112
Po 120	128.121	128	2000	FWD (vlt)	0	0	0001.e88a.dff8 128.121

Connecting a VLT Domain

To connect a VLT domain to an attached access device, use the following commands. For more information, refer to [Verifying a VLT Configuration](#).

1. Configure the VLT domain with the same ID in VLT peer 1 and VLT peer 2.


```
VLT DOMAIN
```

```
vlt domain domain id
```

2. **Configure the VLTi between VLT peer 1 and VLT peer 2.** LACP/Static LAG can be configured between the peer units (not shown).

```
CONFIGURATION mode
```

```
interface port-channel port-channel id
```

 **NOTE:** To benefit from the protocol negotiations, Dell Networking recommends configuring VLTs used as facing hosts/switches with LACP. Both peers must use the same port channel ID.

3. Configure the peer-link port-channel in the VLT domains of each peer unit.

```
INTERFACE PORTCHANNEL mode
```

```
channel-member
```

4. **Configure the backup link between the VLT peer units.** Configure the peer 2 management ip/ interface ip for which connectivity is present in VLT peer 1.

```
EXEC Privilege mode
```

```
show running-config vlt
```

5. Configure the peer 1 management ip/ interface ip for which connectivity is present in VLT peer 1.

```
EXEC mode or EXEC Privilege mode
```

```
show interfaces interface
```

6. **Configure the VLT links between VLT peer 1 and VLT peer 2 to the top of rack unit.** Configure the static LAG/LACP between ports connected from VLT peer 1 and VLT peer 2 to the top of rack unit.

```
EXEC Privilege mode
```

```
show running-config entity
```

7. Configure the VLT peer link port channel id in VLT peer 1 and VLT peer 2.

```
EXEC mode or EXEC Privilege mode
```

```
show interfaces interface
```

8. In the top of rack unit, configure LACP in the physical ports.

```
EXEC Privilege mode
```

```
show running-config entity
```

9. Verify VLT is running.

```
EXEC mode
```

```
show vlt brief
```

```
show vlt detail
```

10. Verify the VLT LAG is running in both VLT peer units.

```
EXEC mode or EXEC Privilege
```

```
show interfaces interface
```

Configure the VLT domain with the same ID in VLT peer 1 and VLT peer 2

Configure the VLTi between VLT peer 1 and VLT peer 2

Configure the backup link between the VLT peer units.


Configure the VLT links between VLT peer 1 and VLT peer 2 to the top of rack unit.

In the ToR unit, configure LACP on the physical ports

Verify VLT is up. Verify that the VLTi (ICL) link, backup link connectivity (heartbeat status) and VLT peer link (peer chassis) are all up.

Verify the VLT LAG is up in both VLT peer units.

In the following sample VLT configuration steps, VLT peer 1 is MXL-2, VLT peer 2 is MXL-4, and the ToR is MXL-1.

 **NOTE:** If you use a third-party ToR unit, Dell Networking recommends using static LAGs with VLT peers to avoid potential problems if the VLT peers are rebooted.

```
mxl-2(conf)#vlt domain 5
mxl-2(conf-vlt-domain)#
```

```
mxl-4(conf)#vlt domain 5
mxl-4(conf-vlt-domain)#
```

1. LACP/Static LAG can be configured between the peer units (not shown).
2. Configure the peer-link port-channel in the VLT domains of each peer unit.

```
mxl-2(conf)#interface port-channel 1
mxl-2(conf-if-po-1)#channel-member TenGigabitEthernet 0/4-7
```

```
mxl-2(conf)#no shutdown
mxl-4(conf)#interface port-channel 1
mxl-4(conf-if-po-1)#channel-member TenGigabitEthernet 0/4-7
mxl-4(conf)#no shutdown
```

1. Configure the peer 2 management ip/ interface ip for which connectivity is present in VLT peer 1.
2. Configure the peer 1 management ip/ interface ip for which connectivity is present in VLT peer 2.

```
mxl-2#show running-config vlt
!
vlt domain 5
  peer-link port-channel 1
  back-up destination 10.11.206.58
mxl-2#
```

```
mxl-2#show interfaces managementethernet 0/0
Internet address is 10.11.206.43/16
```

```
mxl-4#show running-config vlt
!
vlt domain 5
  peer-link port-channel 1
  back-up destination 10.11.206.43
mxl-4#
mxl-4#show running-config interface managementethernet 0/0
ip address 10.11.206.58/16
  no shutdown
```

In the following example, port Te 0/40 in VLT peer 1 is connected to Te 0/48 of TOR and port Te 0/18 in VLT peer 2 is connected to Te 0/50 of TOR.

1. Configure the static LAG/LACP between ports connected from VLT peer 1 and VLT peer 2 to the top of rack unit.
2. Configure the VLT peer link port channel id in VLT peer 1 and VLT peer 2.
3. In the top of rack unit, configure LACP in the physical ports (shown for VLT peer 1 only. Repeat steps for VLT peer 2. The highlighted vlt-peer-lag port-channel 2 indicates that port-channel 2 is the port-channel id configured in VLT peer 2).

```

mxl-2#show running-config interface tengigabitethernet 0/40
!
interface TenGigabitEthernet 0/40
  no ip address
!
port-channel-protocol LACP
  port-channel 2 mode active
  no shutdown
mxl-2#
configuring VLT peer lag in VLT
mxl-2#show running-config interface port-channel 2
!
  interface Port-channel 2
    no ip address
    switchport
    vlt-peer-lag port-channel 2
    no shutdown
mxl-2#
mxl-2#show interfaces port-channel 2 brief
Codes: L - LACP Port-channel

   LAG Mode Status Uptime   Ports
L 2   L2L3 up      03:33:14 Te 0/40 (Up)
mxl-2#

```

```

mxl-4#show running-config interface tengigabitethernet 0/40
!
interface TenGigabitEthernet 0/40
  no ip address
!
port-channel-protocol LACP
  port-channel 2 mode active
  no shutdown
mxl-4#
configuring VLT peer lag in VLT
mxl-4#show running-config interface port-channel 2
!
  interface Port-channel 2
    no ip address
    switchport
    vlt-peer-lag port-channel 2
    no shutdown
mxl-4#
mxl-4#show interfaces port-channel 2 brief
Codes: L - LACP Port-channel

   LAG Mode Status Uptime   Ports
L 2   L2L3 up      03:33:14 Te 0/40 (Up)
mxl-4#

```

```

mxl-1#show running-config interface tengigabitethernet 0/48
!
interface TenGigabitEthernet 0/48
  no ip address
!
  port-channel-protocol LACP
  port-channel 100 mode active
mxl-1#show running-config interface tengigabitethernet 0/50

```

```

!
interface TenGigabitEthernet 0/50
  no ip address
!
port-channel-protocol LACP
port-channel 100 mode active
no shutdown
mxl-1#
mxl-1#show running-config interface port-channel 100
!
interface Port-channel 100
  no ip address
  switchport
  no shutdown
mxl-1#
mxl-1#show interfaces port-channel 100 brief
Codes: L - LACP Port-channel
      LAG Mode Status Uptime  Ports
L 100 L2  up      03:33:48 Te 0/48 (Up)
Te 0/50 (Up)
mxl-1#

Dell#show vlt br
VLT Domain Brief
-----
Domain ID           : 1
Role                : Secondary
Role Priority       : 32768
ICL Link Status    : Up
HeartBeat Status   : Up
VLT Peer Status    : Up
Version            : 6(3)
Local System MAC address : 00:01:e8:8a:e9:91
Remote System MAC address : 00:01:e8:8a:e9:76
Remote system version : 6(3)
Delay-Restore timer : 90 seconds

Delay-Restore Abort Threshold : 60 seconds
Peer-Routing                  : Disabled
Peer-Routing-Timeout timer   : 0 seconds
Multicast peer-routing timeout : 150 seconds
Dell#

Dell#FTOS(conf-if-vl-100)#show vlt detail
Local LAG Id Peer LAG Id Local Status Peer Status Active VLANs
-----
10          10          UP          UP          100, 200, 300, 400,

mxl-2#show interfaces port-channel 2 brief
Codes: L - LACP Port-channel

      LAG Mode Status Uptime  Ports
L 2    L2L3 up      03:43:24 Te 0/40 (Up)
mxl-2#
mxl-4#show interfaces port-channel 2 brief
Codes: L - LACP Port-channel

      LAG Mode Status Uptime  Ports
L 2    L2L3 up      03:33:31 Te 0/18 (Up)
mxl-4#

```

PVST+ Configuration

PVST+ is supported in a VLT domain.

Before you configure VLT on peer switches, configure PVST+ in the network. PVST+ is required for initial loop prevention during the VLT startup phase. You may also use PVST+ for loop prevention in the network outside of the VLT port channel. For information on PVST+, refer to [Per-VLAN Spanning Tree Plus \(PVST+\)](#).

Run PVST+ on both VLT peer switches. PVST+ instance will be created for every VLAN configured in the system. PVST+ instances running in the Primary Peer will control the VLT-LAGs on both Primary and Secondary peers. Only the Primary VLT switch determines the PVST+ roles and states on VLT ports and ensures that the VLT interconnect link is never blocked. PVST+ instance in Primary peer will send the role/state of VLT-LAGs for all VLANs to the Secondary peer. Secondary peer will use this information to program the hardware. PVST+ instance running in Secondary peer will not control the VLT-LAGs.

Dell Networking recommends configuring the primary VLT peer as the primary root device for all the configured PVST+ Instances and configuring the secondary VLT peer as the secondary root device for all the configured PVST+ Instances.

Sample PVST+ Configuration

The following examples show the PVST+ configuration that you must perform on each peer switch to prevent forwarding loops.

Configure PVST+ on VLT Peers to Prevent Forwarding Loops (VLT Peer 1)

```
Dell_VLTpeer1(conf)#protocol spanning-tree pvst
Dell_VLTpeer1(conf-pvst)#no disable
Dell_VLTpeer1(conf-pvst)#vlan 1000 bridge-priority 0
```

Configure PVST+ on VLT Peers to Prevent Forwarding Loops (VLT Peer 2)

```
Dell_VLTpeer2(conf)#protocol spanning-tree pvst
Dell_VLTpeer2(conf-pvst)#no disable
Dell_VLTpeer2(conf-pvst)#vlan 1000 bridge-priority 4096
```

Configure both ends of the VLT interconnect trunk with identical PVST+ configurations. When you enable VLT, the `show spanning-tree pvst brief` command output displays VLT information (refer to [Verifying a VLT Configuration](#)).

```
Dell#show spanning-tree pvst vlan 1000 brief
VLAN 1000
Executing IEEE compatible Spanning Tree Protocol
Root ID      Priority 0, Address 90b1.1cf4.9b79
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID    Priority 0, Address 90b1.1cf4.9b79
We are the root of Vlan 1000
Configured hello time 2, max age 20, forward delay 15
```

```
Interface
Designated
  Name      PortID  Prio Cost  Sts          Cost      Bridge ID
PortID
-----
-----
```



```

Po 1      128.2    128  188    FWD (vltI)  0    0    90b1.1cf4.9b79
128.2
Po 2      128.3    128  2000   FWD (vlt)   0    0    90b1.1cf4.9b79
128.3
Te 0/100  128.230   128  2000   FWD      0    0    90b1.1cf4.9b79
128.230
Te 0/103  128.233   128  2000   FWD      0    0    90b1.1cf4.9b79
128.233

Interface
Name      Role    PortID  Prio  Cost    Sts      Cost    Link-type Edge
-----
Po 1      Desg   128.2   128   188     FWD      0        (vltI) P2P  No
Po 2      Desg   128.3   128   2000    FWD      0        (vlt) P2P  No
Te 0/100  Desg   128.230 128   2000    FWD      0        P2P      Yes
Te 0/103  Desg   128.233 128   2000    FWD      0        P2P      No
Dell#

```

mVLT Configuration Example

The following example demonstrates the steps to configure multi-domain VLT (mVLT) in a network.

In this example there are two domains being configured. Domain 1 consists of Peer 1 and Peer 2; Domain 2 consists of Peer 3 and Peer 4 as shown.

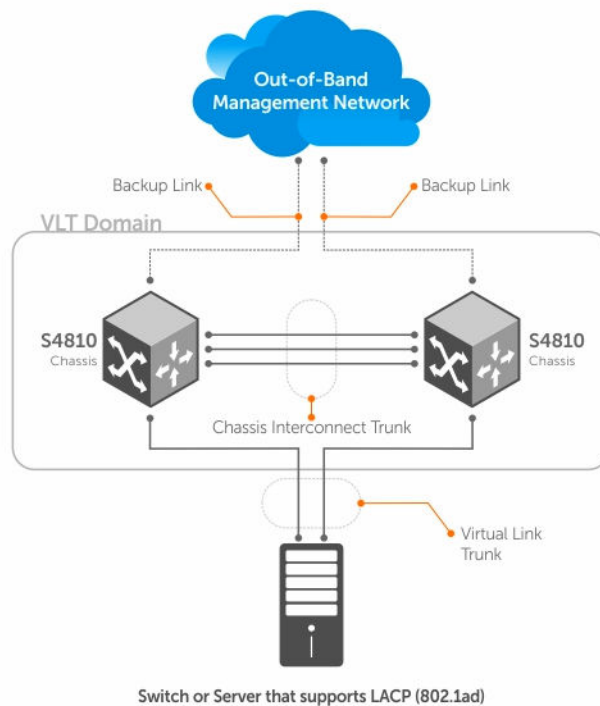


Figure 134. mVLT Configuration Example

In Domain 1, configure Peer 1 first, then configure Peer 2. When that is complete, perform the same steps for the peer nodes in Domain 2. The interface used in this example is TenGigabitEthernet.

In Domain 1, configure the VLT domain and VLTi on Peer 1

```
Domain_1_Peer1#configure
Domain_1_Peer1(conf)#interface port-channel 1
Domain_1_Peer1(conf-if-po-1)#channel-member TenGigabitEthernet 0/8-9
Domain_1_Peer1#no shutdown

Domain_1_Peer1(conf)#vlt domain 1000
Domain_1_Peer1(conf-vlt-domain)#peer-link port-channel 1
Domain_1_Peer1(conf-vlt-domain)#back-up destination 10.16.130.11
Domain_1_Peer1(conf-vlt-domain)#system-mac mac-address 00:0a:00:0a:00:0a
Domain_1_Peer1(conf-vlt-domain)#unit-id 0
```

Configure mVLT on Peer 1

```
Domain_1_Peer1(conf)#interface port-channel 100
Domain_1_Peer1(conf-if-po-100)#switchport
Domain_1_Peer1(conf-if-po-100)#vlt-peer-lag port-channel 100
Domain_1_Peer1(conf-if-po-100)#no shutdown
```

Add links to the mVLT port-channel on Peer 1

```
Domain_1_Peer1(conf)#interface range tengigabitethernet 0/16 - 17
Domain_1_Peer1(conf-if-range-te-0/16-17)#port-channel-protocol LACP
Domain_1_Peer1(conf-if-range-te-0/16-17)#port-channel 100 mode active
Domain_1_Peer1(conf-if-range-te-0/16-17)#no shutdown
```

Next, configure the VLT domain and VLTi on Peer 2

```
Domain_1_Peer2#configure
Domain_1_Peer2(conf)#interface port-channel 1
Domain_1_Peer2(conf-if-po-1)#channel-member TenGigabitEthernet 0/8-9
Domain_1_Peer2#no shutdown

Domain_1_Peer2(conf)#vlt domain 200
Domain_1_Peer2(conf-vlt-domain)#peer-link port-channel 1
Domain_1_Peer2(conf-vlt-domain)#back-up destination 10.16.130.12
Domain_1_Peer2(conf-vlt-domain)#system-mac mac-address 00:0a:00:0a:00:0a
Domain_1_Peer2(conf-vlt-domain)#unit-id 1
```

Configure mVLT on Peer 2

```
Domain_1_Peer2(conf)#interface port-channel 100
Domain_1_Peer2(conf-if-po-100)#switchport
Domain_1_Peer2(conf-if-po-100)#vlt-peer-lag port-channel 100
Domain_1_Peer2(conf-if-po-100)#no shutdown
```

Add links to the mVLT port-channel on Peer 2

```
Domain_1_Peer2(conf)#interface range tengigabitethernet 0/28 - 29
Domain_1_Peer2(conf-if-range-te-0/16-17)#port-channel-protocol LACP
Domain_1_Peer2(conf-if-range-te-0/16-17)#port-channel 100 mode active
Domain_1_Peer2(conf-if-range-te-0/16-17)#no shutdown
```

In Domain 2, configure the VLT domain and VLTi on Peer 3

```
Domain_2_Peer3#configure
Domain_2_Peer3(conf)#interface port-channel 1
Domain_2_Peer3(conf-if-po-1)#channel-member TenGigabitEthernet 0/8-9
```

```
Domain_1_Peer3#no shutdown
Domain_2_Peer3(conf)#vlt domain 200
Domain_2_Peer3(conf-vlt-domain)#peer-link port-channel 1
Domain_2_Peer3(conf-vlt-domain)#back-up destination 10.18.130.11
Domain_2_Peer3(conf-vlt-domain)#system-mac mac-address 00:0b:00:0b:00:0b
Domain_2_Peer3(conf-vlt-domain)#unit-id 0
```

Configure mVLT on Peer 3

```
Domain_2_Peer3(conf)#interface port-channel 100
Domain_2_Peer3(conf-if-po-100)#switchport
Domain_2_Peer3(conf-if-po-100)#vlt-peer-lag port-channel 100
Domain_2_Peer3(conf-if-po-100)#no shutdown
```

Add links to the mVLT port-channel on Peer 3

```
Domain_2_Peer3(conf)#interface range tengigabitethernet 0/19 - 20
Domain_2_Peer3(conf-if-range-te-0/16-17)#port-channel-protocol LACP
Domain_2_Peer3(conf-if-range-te-0/16-17)#port-channel 100 mode active
Domain_2_Peer3(conf-if-range-te-0/16-17)#no shutdown
```

Configure the VLT domain and VLTi on Peer 4

```
Domain_2_Peer4#configure
Domain_2_Peer4(conf)#interface port-channel 1
Domain_2_Peer4(conf-if-po-1)#channel-member TenGigabitEthernet 0/8-9
Domain_1_Peer4#no shutdown
```

```
Domain_2_Peer4(conf)#vlt domain 200
Domain_2_Peer4(conf-vlt-domain)#peer-link port-channel 1
Domain_2_Peer4(conf-vlt-domain)#back-up destination 10.18.130.12
Domain_2_Peer4(conf-vlt-domain)#system-mac mac-address 00:0b:00:0b:00:0b
Domain_2_Peer4(conf-vlt-domain)#unit-id 1
```

Configure mVLT on Peer 4

```
Domain_2_Peer4(conf)#interface port-channel 100
Domain_2_Peer4(conf-if-po-100)#switchport
Domain_2_Peer4(conf-if-po-100)#vlt-peer-lag port-channel 100
Domain_2_Peer4(conf-if-po-100)#no shutdown
```

Add links to the mVLT port-channel on Peer 4

```
Domain_2_Peer4(conf)#interface range tengigabitethernet 0/31 - 32
Domain_2_Peer4(conf-if-range-te-0/16-17)#port-channel-protocol LACP
Domain_2_Peer4(conf-if-range-te-0/16-17)#port-channel 100 mode active
Domain_2_Peer4(conf-if-range-te-0/16-17)#no shutdown
```

PIM-Sparse Mode Configuration Example

The following sample configuration shows how to configure the PIM Sparse mode designated router functionality on the VLT domain with two VLT port-channels that are members of VLAN 4001.

For more information, refer to [PIM-Sparse Mode Support on VLT](#).

Example of Configuring PIM-Sparse Mode

Enable PIM Multicast Routing on the VLT node globally.

```
VLT_Peer1(conf)#ip multicast-routing
```

Enable PIM on the VLT port VLANs.

```
VLT_Peer1(conf)#interface vlan 4001
VLT_Peer1(conf-if-vl-4001)#ip address 140.0.0.1/24
VLT_Peer1(conf-if-vl-4001)#ip pim sparse-mode
VLT_Peer1(conf-if-vl-4001)#tagged port-channel 101
VLT_Peer1(conf-if-vl-4001)#tagged port-channel 102
VLT_Peer1(conf-if-vl-4001)#no shutdown
VLT_Peer1(conf-if-vl-4001)#exit
```

Configure the VLTi port as a static multicast router port for the VLAN.

```
VLT_Peer1(conf)#interface vlan 4001
VLT_Peer1(conf-if-vl-4001)#ip igmp snooping mrouter interface port-channel 128
VLT_Peer1(conf-if-vl-4001)#exit
VLT_Peer1(conf)#end
```

Repeat these steps on VLT Peer Node 2.

```
VLT_Peer2(conf)#ip multicast-routing

VLT_Peer2(conf)#interface vlan 4001
VLT_Peer2(conf-if-vl-4001)#ip address 140.0.0.2/24
VLT_Peer2(conf-if-vl-4001)#ip pim sparse-mode
VLT_Peer2(conf-if-vl-4001)#tagged port-channel 101
VLT_Peer2(conf-if-vl-4001)#tagged port-channel 102
VLT_Peer2(conf-if-vl-4001)#no shutdown

VLT_Peer2(conf-if-vl-4001)#ip igmp snooping mrouter interface port-channel 128
VLT_Peer2(conf-if-vl-4001)#exit
VLT_Peer2(conf)#end
```

Additional VLT Sample Configurations

To configure VLT, configure a backup link and interconnect trunk, create a VLT domain, configure a backup link and interconnect trunk, and connect the peer switches in a VLT domain to an attached access device (switch or server).

Review the following examples of VLT configurations.

Configuring Virtual Link Trunking (VLT Peer 1)

Enable VLT and create a VLT domain with a backup-link and interconnect trunk (VLTi).

```
Dell_VLTpeer1(conf)#vlt domain 999
Dell_VLTpeer1(conf-vlt-domain)#peer-link port-channel 100
Dell_VLTpeer1(conf-vlt-domain)#back-up destination 10.11.206.35
Dell_VLTpeer1(conf-vlt-domain)#exit
```

Configure the backup link.

```
Dell_VLTpeer1(conf)#interface ManagementEthernet 0/0
Dell_VLTpeer1(conf-if-ma-0/0)#ip address 10.11.206.23/
Dell_VLTpeer1(conf-if-ma-0/0)#no shutdown
Dell_VLTpeer1(conf-if-ma-0/0)#exit
```

Configure the VLT interconnect (VLTi).

```
Dell_VLTpeer1(conf)#interface port-channel 100
Dell_VLTpeer1(conf-if-po-100)#no ip address
Dell_VLTpeer1(conf-if-po-100)#channel-member fortyGigE 0/56,60
Dell_VLTpeer1(conf-if-po-100)#no shutdown
Dell_VLTpeer1(conf-if-po-100)#exit
```

Configure the port channel to an attached device.

```
Dell_VLTpeer1(conf)#interface port-channel 110
Dell_VLTpeer1(conf-if-po-110)#no ip address
Dell_VLTpeer1(conf-if-po-110)#switchport
Dell_VLTpeer1(conf-if-po-110)#channel-member fortyGigE 0/52
Dell_VLTpeer1(conf-if-po-110)#no shutdown
Dell_VLTpeer1(conf-if-po-110)#vlt-peer-lag port-channel 110
Dell_VLTpeer1(conf-if-po-110)#end
```

Verify that the port channels used in the VLT domain are assigned to the same VLAN.

```
Dell_VLTpeer1# show vlan id 10
Codes: * - Default VLAN, G - GVRP VLANs, P - Primary, C - Community, I -
Isolated
Q: U - Untagged, T - Tagged
   x - Dot1x untagged, X - Dot1x tagged
   G - GVRP tagged, M - Vlan-stack, H - Hyperpull tagged

  NUM Status Description Q Ports
  10  Active                U Po110(Fo 0/52)
                        T Po100(Fo 0/56,60)
```

Configuring Virtual Link Trunking (VLT Peer 2)

Enable VLT and create a VLT domain with a backup-link VLT interconnect (VLTi).

```
Dell_VLTpeer2(conf)#vlt domain 999
Dell_VLTpeer2(conf-vlt-domain)#peer-link port-channel 100
Dell_VLTpeer2(conf-vlt-domain)#back-up destination 10.11.206.23
Dell_VLTpeer2(conf-vlt-domain)#exit
```

Configure the backup link.

```
Dell_VLTpeer2(conf)#interface ManagementEthernet 0/0
Dell_VLTpeer2(conf-if-ma-0/0)#ip address 10.11.206.35/
Dell_VLTpeer2(conf-if-ma-0/0)#no shutdown
Dell_VLTpeer2(conf-if-ma-0/0)#exit
```

Configure the VLT interconnect (VLTi).

```
Dell_VLTpeer2(conf)#interface port-channel 100
Dell_VLTpeer2(conf-if-po-100)#no ip address
Dell_VLTpeer2(conf-if-po-100)#channel-member fortyGigE 0/46,50
Dell_VLTpeer2(conf-if-po-100)#no shutdown
Dell_VLTpeer2(conf-if-po-100)#exit
```

Configure the port channel to an attached device.

```
Dell_VLTpeer2(conf)#interface port-channel 110
Dell_VLTpeer2(conf-if-po-110)#no ip address
Dell_VLTpeer2(conf-if-po-110)#switchport
Dell_VLTpeer2(conf-if-po-110)#channel-member fortyGigE 0/48
Dell_VLTpeer2(conf-if-po-110)#no shutdown
Dell_VLTpeer2(conf-if-po-110)#vlt-peer-lag port-channel 110
Dell_VLTpeer2(conf-if-po-110)#end
```

Verify that the port channels used in the VLT domain are assigned to the same VLAN.

```
Dell_VLTpeer2# show vlan id 10
Codes: * - Default VLAN, G - GVRP VLANs, P - Primary, C - Community, I -
Isolated
Q: U - Untagged, T - Tagged
   x - Dot1x untagged, X - Dot1x tagged
   G - GVRP tagged, M - Vlan-stack, H - Hyperpull tagged

   NUM Status Description Q Ports
   10 Active                U Po110(Fo 0/48)
                           T Po100(Fo 0/46,50)
```

Verifying a Port-Channel Connection to a VLT Domain (From an Attached Access Switch)

On an access device, verify the port-channel connection to a VLT domain.

```
Dell_TORswitch(conf)# show running-config interface port-channel 11
!
interface Port-channel 11
no ip address
switchport
channel-member fortyGigE 1/18,22
no shutdown
```

Troubleshooting VLT

To help troubleshoot different VLT issues that may occur, use the following information.


 **NOTE:** For information on VLT Failure mode timing and its impact, contact your Dell Networking representative.

Table 69. Troubleshooting VLT

Description	Behavior at Peer Up	Behavior During Run Time	Action to Take
Bandwidth monitoring	A syslog error message and an SNMP trap is generated when the VLTi bandwidth usage goes above the 80% threshold and when it drops below 80%.	A syslog error message and an SNMP trap is generated when the VLTi bandwidth usage goes above its threshold.	Depending on the traffic that is received, the traffic can be offloaded inVLTi.
Domain ID mismatch	The VLT peer does not boot up. The VLTi is forced to a down state. A syslog error message and an SNMP trap are generated.	The VLT peer does not boot up. The VLTi is forced to a down state. A syslog error message and an SNMP trap are generated.	Verify the domain ID matches on both VLT peers.
Dell Networking Version mismatch	A syslog error message is generated.	A syslog error message is generated.	Follow the correct upgrade procedure for the unit with the mismatched Dell Networking version.
Remote VLT port channel status	N/A	N/A	Use the <code>show vlt detail</code> and <code>show vlt brief</code> commands to view the VLT port channel status information.
Spanning tree mismatch at global level	All VLT port channels go down on both VLT peers. A syslog error message is generated.	No traffic is passed on the port channels. A one-time informational syslog message is generated.	During run time, a loop may occur as long as the mismatch lasts. To resolve, enable RSTP on both VLT peers.
Spanning tree mismatch at port level	A syslog error message is generated.	A one-time informational syslog message is generated.	Correct the spanning tree configuration on the ports.
System MAC mismatch	A syslog error message and an SNMP trap are generated.	A syslog error message and an SNMP trap are generated.	Verify that the unit ID of VLT peers is not the same on both units and that the MAC address is the same on both units.
Unit ID mismatch	The VLT peer does not boot up. The VLTi is forced to a down state.	The VLT peer does not boot up. The VLTi is forced to a down state.	Verify the unit ID is correct on both VLT peers. Unit ID numbers must be sequential on peer units; for example,

Description	Behavior at Peer Up	Behavior During Run Time	Action to Take
	A syslog error message is generated.	A syslog error message is generated.	if Peer 1 is unit ID "0", Peer 2 unit ID must be "1".
Version ID mismatch	A syslog error message and an SNMP trap are generated.	A syslog error message and an SNMP trap are generated.	Verify the Dell Networking software versions on the VLT peers is compatible. For more information, refer to the <i>Release Notes</i> for this release.
VLT LAG ID is not configured on one VLT peer	A syslog error message is generated. The peer with the VLT configured remains active.	A syslog error message is generated. The peer with the VLT configured remains active.	Verify the VLT LAG ID is configured correctly on both VLT peers.
VLT LAG ID mismatch	The VLT port channel is brought down. A syslog error message is generated.	The VLT port channel is brought down. A syslog error message is generated.	Perform a mismatch check after the VLT peer is established.
VLT LAG VLAN mismatch	A syslog error message is generated.	A syslog error message is generated.	Verify that the VLAN configuration is same for the VLT lags on both peers.

Specifying VLT Nodes in a PVLAN

VLT enables redundancy without the implementation of Spanning Tree Protocol (STP), and provides a loop-free network with optimal bandwidth utilization.

Because the VLT LAG interfaces are terminated on two different nodes, PVLAN configuration of VLT VLANs and VLT LAGs are symmetrical and identical on both the VLT peers. PVLANS provide Layer 2 isolation between ports within the same VLAN. A PVLAN partitions a traditional VLAN into sub-domains identified by a primary and secondary VLAN pair. With VLT being a Layer 2 redundancy mechanism, support for configuration of VLT nodes in a PVLAN enables Layer 2 security functionalities. To achieve maximum VLT resiliency, you should configure the PVLAN IDs and mappings to be identical on both the VLT peer nodes.

The association of PVLAN with the VLT LAG must also be identical. After the VLT LAG is configured to be a member of either the primary or secondary PVLAN (which is associated with the primary), ICL becomes an automatic member of that PVLAN on both switches. This association helps the PVLAN data flow received on one VLT peer for a VLT LAG to be transmitted on that VLT LAG from the peer.

You can associate either a VLT VLAN or a VLT LAG to a PVLAN. First configure the VLT interconnect (VLTi) or a VLT LAG by using the `peer-link port-channel id-number` command or the VLT VLAN by using

the `peer-link port-channel id-number peer-down-vlan vlan interface number` command and the `switchport` command. After you specify the VLTi link and VLT LAGs, you can associate the same port channel or LAG bundle that is a part of a VLT to a PVLAN by using the `interface interface` and `switchport mode private-vlan` commands.

When a VLTi port in trunk mode is a member of symmetric VLT PVLANs, the PVLAN packets are forwarded only if the PVLAN settings of both the VLT nodes are identical. You can configure the VLTi in trunk mode to be a member of non-VLT PVLANs if the VLTi is configured on both the peers. MAC address synchronization is performed for VLT PVLANs across peers in a VLT domain.

Keep the following points in mind when you configure VLT nodes in a PVLAN:

- Configure the VLTi link to be in trunk mode. Do not configure the VLTi link to be in access or promiscuous mode.
- You can configure a VLT LAG or port channel to be in trunk, access, or promiscuous port modes when you include the VLT LAG in a PVLAN. The VLT LAG settings must be the same on both the peers. If you configure a VLT LAG as a trunk port, you can associate that LAG to be a member of a normal VLAN or a PVLAN. If you configure a VLT LAG to be a promiscuous port, you can configure that LAG to be a member of PVLAN only. If you configure a VLT LAG to be in access port mode, you can add that LAG to be a member of the secondary VLAN only.
- ARP entries are synchronized even when a mismatch occurs in the PVLAN mode of a VLT LAG.

Any VLAN that contains at least one VLT port as a member is treated as a VLT VLAN. You can configure a VLT VLAN to be a primary, secondary, or a normal VLAN. However, the VLT VLAN configuration must be symmetrical across peers. If the VLT LAG is tagged to any one of the primary or secondary VLANs of a PVLAN, then both the primary and secondary VLANs are considered as VLT VLANs.

If you add an ICL or VLTi link as a member of a primary VLAN, the ICL becomes a part of the primary VLAN and its associated secondary VLANs, similar to the behavior for normal trunk ports. VLAN parity is not validated if you associate an ICL to a PVLAN. Similarly, if you dissociate an ICL from a PVLAN, although the PVLAN parity exists, ICL is removed from that PVLAN.

Association of VLTi as a Member of a PVLAN

If a VLAN is configured as a non-VLT VLAN on both the peers, the VLTi link is made a member of that VLAN if the VLTi link is configured as a PVLAN or normal VLAN on both the peers. If a PVLAN is configured as a VLT VLAN on one peer and a non-VLT VLAN on another peer, the VLTi is added as a member of that VLAN by verifying the PVLAN parity on both the peers. In such a case, if a PVLAN is present as a VLT PVLAN on at least one of the peers, then symmetric configuration of the PVLAN is validated to cause the VLTi to be a member of that VLAN. Whenever a change in the VLAN mode on one of the peers occurs, the information is synchronized with the other peer and VLTi is either added or removed from the VLAN based on the validation of the VLAN parity.

For VLT VLANs, the association between primary VLAN and secondary VLANs is examined on both the peers. Only if the association is identical on both the peers, VLTi is configured as a member of those VLANs. This behavior is because of security functionalities in a PVLAN. For example, if a VLAN is a primary VLT VLAN on one peer and not a primary VLT VLAN on the other peer, VLTi is not made a part of that VLAN.

MAC Synchronization for VLT Nodes in a PVLAN

For the MAC addresses that are learned on non-VLT ports, MAC address synchronization is performed with the other peer if the VLTi (ICL) link is part of the same VLAN as the non-VLT port. For MAC addresses that are learned on VLT ports, the VLT LAG mode of operation and the primary to secondary association of the VLT nodes is determined on both the VLT peers. MAC synchronization is performed for the VLT LAGs only if the VLT LAG and primary-secondary VLT peer mapping are symmetrical.

The PVLAN mode of VLT LAGs on one peer is validated against the PVLAN mode of VLT LAGs on the other peer. MAC addresses that are learned on that VLT LAG are synchronized between the peers only if the PVLAN mode on both the peers is identical. For example, if the MAC address is learned on a VLT LAG and the VLAN is a primary VLT VLAN on one peer and not a primary VLT VLAN on the other peer, MAC synchronization does not occur.

Whenever a change occurs in the VLAN mode of one of the peers, this modification is synchronized with the other peers. Depending on the validation mechanism that is initiated for MAC synchronization of VLT peers, MAC addresses learned on a particular VLAN are either synchronized with the other peers, or MAC addresses synchronized from the other peers on the same VLAN are deleted. This method of processing occurs when the PVLAN mode of VLT LAGs is modified.

Because the VLTi link is only a member of symmetric VLT PVLANS, MAC synchronization takes place directly based on the membership of the VLTi link in a VLAN and the VLT LAG mode.

PVLAN Operations When One VLT Peer is Down

When a VLT port moves to the Admin or Operationally Down state on only one of the VLT nodes, the VLT Lag is still considered to be up. All the PVLAN MAC entries that correspond to the operationally down VLT LAG are maintained as synchronized entries in the device. These MAC entries are removed when the peer VLT LAG also becomes inactive or a change in PVLAN configuration occurs.

PVLAN Operations When a VLT Peer is Restarted

When the VLT peer node is rebooted, the VLAN membership of the VLTi link is preserved and when the peer node comes back online, a verification is performed with the newly received PVLAN configuration from the peer. If any differences are identified, the VLTi link is either added or removed from the VLAN. When the peer node restarts and returns online, all the PVLAN configurations are exchanged across the peers. Based on the information received from the peer, a bulk synchronization of MAC addresses that belong to spanned PVLANS is performed.

During the booting phase or when the ICL link attempts to come up, a system logging message is recorded if VLT PVLAN mismatches, PVLAN mode mismatches, PVLAN association mismatches, or PVLAN port mode mismatches occur. Also, you can view these discrepancies if any occur by using the `show vlt mismatch` command.

Interoperation of VLT Nodes in a PVLAN with ARP Requests

When an ARP request is received, and the following conditions are applicable, the IP stack performs certain operations.

- The VLAN on which the ARP request is received is a secondary VLAN (community or isolated VLAN).

- Layer 3 communication between secondary VLANs in a private VLAN is enabled by using the `ip local-proxy-arp` command in INTERFACE VLAN configuration mode.
- The ARP request is not received on the ICL

Under such conditions, the IP stack performs the following operations:

- The ARP reply is sent with the MAC address of the primary VLAN.
- The ARP request packet originates on the primary VLAN for the intended destination IP address.

The ARP request received on ICLs are not proxied, even if they are received with a secondary VLAN tag. This behavior change occurs because the node from which the ARP request was forwarded would have replied with its MAC address, and the current node discards the ARP request.

Scenarios for VLAN Membership and MAC Synchronization With VLT Nodes in PVLAN

The following table illustrates the association of the VLTi link and PVLANS, and the MAC synchronization of VLT nodes in a PVLAN (for various modes of operations of the VLT peers):

Table 70. VLAN Membership and MAC Synchronization With VLT Nodes in PVLAN

VLT LAG Mode		PVLAN Mode of VLT VLAN		ICL VLAN Membership	Mac Synchronization
Peer1	Peer2	Peer1	Peer2		
Trunk	Trunk	Primary	Primary	Yes	Yes
Trunk	Trunk	Primary	Normal	No	No
Trunk	Trunk	Normal	Normal	Yes	Yes
Promiscuous	Trunk	Primary	Primary	Yes	No
Trunk	Access	Primary	Secondary	No	No
Promiscuous	Promiscuous	Primary	Primary	Yes	Yes
Promiscuous	Access	Primary	Secondary	No	No
Promiscuous	Promiscuous	Primary	Primary	Yes	Yes
		- Secondary (Community)	- Secondary (Isolated)	No	No
Access	Access	Secondary (Community)	Secondary (Isolated)	No	No
		• Primary X	• Primary X	Yes	Yes
Promiscuous	Promiscuous	Primary	Primary	Yes	Yes

VLT LAG Mode		PVLAN Mode of VLT VLAN		ICL VLAN Membership	Mac Synchronization
Peer1	Peer2	Peer1	Peer2		
		- Secondary (Community)	- Secondary (Community)	Yes	Yes
		- Secondary (Isolated)	- Secondary (Isolated)	Yes	Yes
Promiscuous	Trunk	Primary	Normal	No	No
Promiscuous	Trunk	Primary	Primary	Yes	No
Access	Access	Secondary (Community)	Secondary (Community)	Yes	Yes
		- Primary VLAN X	- Primary VLAN X	Yes	Yes
Access	Access	Secondary (Isolated)	Secondary (Isolated)	Yes	Yes
		- Primary VLAN X	- Primary VLAN X	Yes	Yes
Access	Access	Secondary (Isolated)	Secondary (Isolated)	No	No
		- Primary VLAN X	- Primary VLAN Y	No	No
Access	Access	Secondary (Community)	Secondary (Community)	No	No
		- Primary VLAN Y	- Primary VLAN X	No	No
Promiscuous	Access	Primary	Secondary	No	No
Trunk	Access	Primary/Normal	Secondary	No	No

Configuring a VLT VLAN or LAG in a PVLAN


You can configure the VLT peers or nodes in a private VLAN (PVLAN). Because the VLT LAG interfaces are terminated on two different nodes, PVLAN configuration of VLT VLANs and VLT LAGs are symmetrical and identical on both the VLT peers. PVLANS provide Layer 2 isolation between ports within the same VLAN. A PVLAN partitions a traditional VLAN into subdomains identified by a primary and secondary VLAN pair. With VLT being a Layer 2 redundancy feature, support for configuration of VLT nodes in a PVLAN enables Layer 2 security functionalities to be achieved. This section contains the following topics that describe how to configure a VLT VLAN or a VLT LAG (VLTi link) and assign that VLT interface to a PVLAN.

Creating a VLT LAG or a VLT VLAN

1. Configure the port channel for the VLT interconnect on a VLT switch and enter interface configuration mode
CONFIGURATION mode

```
interface port-channel id-number.
```

Enter the same port-channel number configured with the `peer-link port-channel` command.

 **NOTE:** To be included in the VLTi, the port channel must be in Default mode (no `switchport` or VLAN assigned).

2. Remove an IP address from the interface.
INTERFACE PORT-CHANNEL mode

```
no ip address
```

3. Add one or more port interfaces to the port channel.
INTERFACE PORT-CHANNEL mode

```
channel-member interface
```

interface: specify one of the following interface types:

- 1-Gigabit Ethernet: Enter `gigabitethernet slot/port`.
- 10-Gigabit Ethernet: Enter `tengigabitethernet slot/port`.

4. Ensure that the port channel is active.
INTERFACE PORT-CHANNEL mode

```
no shutdown
```

5. To configure the VLT interconnect, repeat Steps 1–4 on the VLT peer switch.
6. Enter VLT-domain configuration mode for a specified VLT domain.
CONFIGURATION mode

```
vlt domain domain-id
```

The range of domain IDs is from 1 to 1000.

7. Enter the port-channel number that acts as the interconnect trunk.
VLT DOMAIN CONFIGURATION mode

```
peer-link port-channel id-number
```

The range is from 1 to 128.

8. (Optional) To configure a VLT LAG, enter the VLAN ID number of the VLAN where the VLT forwards packets received on the VLTi from an adjacent peer that is down.
VLT DOMAIN CONFIGURATION mode

```
peer-link port-channel id-number peer-down-vlan vlan interface number
```

The range is from 1 to 4094.

Associating the VLT LAG or VLT VLAN in a PVLAN

1. Access INTERFACE mode for the port that you want to assign to a PVLAN.
CONFIGURATION mode

```
interface interface
```

2. Enable the port.
INTERFACE mode

```
no shutdown
```

3. Set the port in Layer 2 mode.
INTERFACE mode

```
switchport
```

4. Select the PVLAN mode.
INTERFACE mode

```
switchport mode private-vlan {host | promiscuous | trunk}
```

- `host` (isolated or community VLAN port)
- `promiscuous` (intra-VLAN communication port)
- `trunk` (inter-switch PVLAN hub port)

5. Access INTERFACE VLAN mode for the VLAN to which you want to assign the PVLAN interfaces.
CONFIGURATION mode

```
interface vlan vlan-id
```

6. Enable the VLAN.
INTERFACE VLAN mode

```
no shutdown
```

7. To obtain maximum VLT resiliency, configure the PVLAN IDs and mappings to be identical on both the VLT peer nodes. Set the PVLAN mode of the selected VLAN to primary.

```
INTERFACE VLAN mode
```

```
private-vlan mode primary
```

8. Map secondary VLANs to the selected primary VLAN.
INTERFACE VLAN mode

```
private-vlan mapping secondary-vlan vlan-list
```

The list of secondary VLANs can be:

- Specified in comma-delimited (*VLAN-ID, VLAN-ID*) or hyphenated-range format (*VLAN-ID-VLAN-ID*).
- Specified with this command even before they have been created.
- Amended by specifying the new secondary VLAN to be added to the list.

Proxy ARP Capability on VLT Peer Nodes

A proxy ARP-enabled device answers the ARP requests that are destined for another host or router. The local host forwards the traffic to the proxy ARP-enabled device, which in turn transmits the packets to the destination.

By default, proxy ARP is enabled. To disable proxy ARP, use the `no proxy-arp` command in the interface mode. To re-enable proxy ARP, use the `ip proxy-arp` command in INTERFACE mode. To view if proxy ARP is enabled on the interface, use the `show config` command in INTERFACE mode. If it is not listed in the `show config` command output, it is enabled. Only nondefault information is displayed in the `show config` command output.

ARP proxy operation is performed on the VLT peer node IP address when the peer VLT node is down. The ARP proxy stops working either when the peer routing timer expires or when the peer VLT node goes up. Layer 3 VLT provides a higher resiliency at the Layer 3 forwarding level. VLT peer routing enables you to replace VRRP with routed VLT to route the traffic from Layer 2 access nodes. With proxy ARP, hosts can resolve the MAC address of the VLT node even when VLT node is down.

If the ICL link is down when a VLT node receives an ARP request for the IP address of the VLT peer, owing to LAG-level hashing algorithm in the top-of-rack (TOR) switch, the incorrect VLT node responds to the ARP request with the peer MAC address. Proxy ARP is not performed when the ICL link is up and the ARP request the wrong VLT peer. In this case, ARP requests are tunneled to the VLT peer.

Proxy ARP supported on both VLT interfaces and non-VLT interfaces. Proxy ARP supported on symmetric VLANs only. Proxy ARP is enabled by default. Routing table must be symmetrically configured to support proxy ARP. For example, consider a sample topology in which VLAN 100 is configured on two VLT nodes, node 1 and node 2. ICL link is not configured between the two VLT nodes. Assume that the VLAN 100 IP address in node 1 is 10.1.1.1/24 and VLAN 100 IP address in node 2 is 20.1.1.2/24. In this case, if the ARP request for 20.1.1.1 reaches node 1, node 1 will not perform the ARP request for 20.1.1.2. Proxy ARP is supported only for the IP address belongs to the received interface IP network. Proxy ARP is not supported if the ARP requested IP address is different from the received interface IP subnet. For example, if VLAN 100 and 200 are configured on the VLT peers, and if the VLAN 100 IP address is configured as 10.1.1.0/24 and the VLAN 200 IP address is configured as 20.1.1.0/24, the proxy ARP is not performed if the VLT node receives an ARP request for 20.1.1.0/24 on VLAN 100.

Working of Proxy ARP for VLT Peer Nodes

Proxy ARP is enabled only when peer routing is enabled on both the VLT peers. If peer routing is disabled on one of the VLT peers, proxy ARP is not performed when the ICL link goes down. Proxy ARP is performed only when the VLT peer's MAC address is installed in the database. Proxy ARP is stopped when the VLT peer's MAC address is removed from the ARP database because of the peer routing timer expiry. The source hardware address in the ARP response contains the VLT peer MAC address. Proxy ARP is supported for both unicast and broadcast ARP requests. Control packets, other than ARP requests destined for the VLT peers that reach the undesired and incorrect VLT node, are dropped if the ICL link is down. Further processing is not done on these control packets. The VLT node does not perform any action if it receives gratuitous ARP requests for the VLT peer IP address. Proxy ARP is also supported on

secondary VLANs. When the ICL link or peer is down, and the ARP request for a private VLAN IP address reaches the wrong peer, then the wrong peer responds to the ARP request with the peer MAC address.

The IP address of the VLT node VLAN interface is synchronized with the VLT peer over ICL when the VLT peers are up. Whenever an IP address is added or deleted, this updated information is synchronized with the VLT peer. IP address synchronization occurs regardless of the VLAN administrative state. IP address addition and deletion serve as the trigger events for synchronization. When a VLAN state is down, the VLT peer might perform a proxy ARP operation for the IP addresses of that VLAN interface.

VLT nodes start performing Proxy ARP when the ICL link goes down. When the VLT peer comes up, proxy ARP will be stopped for the peer VLT IP addresses. When the peer node is rebooted, the IP address synchronized with the peer is not flushed. Peer down events cause the proxy ARP to commence.

When a VLT node detects peer up, it will not perform proxy ARP for the peer IP addresses. IP address synchronization occurs again between the VLT peers.

Proxy ARP is enabled only if peer routing is enabled on both the VLT peers. If you disable peer routing by using the `no peer-routing` command in VLT DOMAIN node, a notification is sent to the VLT peer to disable the proxy ARP. If peer routing is disabled when ICL link is down, a notification is not sent to the VLT peer and in such a case, the VLT peer does not disable the proxy ARP operation.

When the VLT domain is removed on one of the VLT nodes, the peer routing configuration removal will be notified to the peer. In this case VLT peer node disables the proxy ARP. When the ICL link is removed on one of the VLT nodes by using the `no peer-link` command, the ICL down event is triggered on the other VLT node, which in turn starts the proxy ARP application. The VLT node, where the ICL link is deleted, flushes the peer IP addresses and does not perform proxy ARP for the additional LAG hashed ARP requests.

Configuring VLAN-Stack over VLT

To configure VLAN-stack over VLT, follow these steps.

1. Configure the VLT LAG as VLAN-stack access or trunk mode on both the peers.

```
INTERFACE PORT-CHANNEL mode
```

```
vlan-stack {access | trunk}
```

2. Configure VLAN as VLAN-stack compatible on both the peers.

```
INTERFACE VLAN mode
```

```
vlan-stack compatible
```

3. Add the VLT LAG as a member to the VLAN-stack on both the peers.

```
INTERFACE VLAN mode
```

```
member port-channel port-channel ID
```

4. Verify the VLAN-stack configurations.

```
EXEC Privilege
```

```
show running-config
```


Sample configuration of VLAN-stack over VLT (Peer 1)

Configure VLT domain

```
Dell(conf)#vlt domain 1
Dell(conf-vlt-domain)#peer-link port-channel 1
Dell(conf-vlt-domain)#back-up destination 10.16.151.116
Dell(conf-vlt-domain)#primary-priority 100
Dell(conf-vlt-domain)#system-mac mac-address 00:00:00:11:11:11
Dell(conf-vlt-domain)#unit-id 0
Dell(conf-vlt-domain)#
```

```
Dell#show running-config vlt
!
vlt domain 1
  peer-link port-channel 1
  back-up destination 10.16.151.116
  primary-priority 100
  system-mac mac-address 00:00:00:11:11:11
  unit-id 0
Dell#
```

Configure VLT LAG as VLAN-Stack Access or Trunk Port

```
Dell(conf)#interface port-channel 10
Dell(conf-if-po-10)#switchport
Dell(conf-if-po-10)#vlt-peer-lag port-channel 10
Dell(conf-if-po-10)#vlan-stack access
Dell(conf-if-po-10)#no shutdown
```

```
Dell#show running-config interface port-channel 10
!
interface Port-channel 10
  no ip address
  switchport
  vlan-stack access
  vlt-peer-lag port-channel 10
  no shutdown
Dell#
```

```
Dell(conf)#interface port-channel 20
Dell(conf-if-po-20)#switchport
Dell(conf-if-po-20)#vlt-peer-lag port-channel 20
Dell(conf-if-po-20)#vlan-stack trunk
Dell(conf-if-po-20)#no shutdown
```

```
Dell#show running-config interface port-channel 20
!
interface Port-channel 20
  no ip address
  switchport
  vlan-stack trunk
  vlt-peer-lag port-channel 20
  no shutdown
Dell#
```

Configure VLAN as VLAN-Stack VLAN and add the VLT LAG as Members to the VLAN

```
Dell(conf)#interface vlan 50
Dell(conf-if-vl-50)#vlan-stack compatible
```

```
Dell(conf-if-vl-50-stack)#member port-channel 10
Dell(conf-if-vl-50-stack)#member port-channel 20
```

```
Dell#show running-config interface vlan 50
!
interface Vlan 50
  vlan-stack compatible
  member Port-channel 10,20
  shutdown
Dell#
```

Verify that the Port Channels used in the VLT Domain are Assigned to the VLAN-Stack VLAN

Sample Configuration of VLAN-Stack Over VLT (Peer 2)

Configure VLT domain

```
Dell(conf)#vlt domain 1
Dell(conf-vlt-domain)#peer-link port-channel 1
Dell(conf-vlt-domain)#back-up destination 10.16.151.115
Dell(conf-vlt-domain)#system-mac mac-address 00:00:00:11:11:11
Dell(conf-vlt-domain)#unit-id 1
Dell(conf-vlt-domain)#
```

```
Dell#show running-config vlt
vlt domain 1
  peer-link port-channel 1
  back-up destination 10.16.151.115
  system-mac mac-address 00:00:00:11:11:11
  unit-id 1
Dell#
```

Configure VLT LAG as VLAN-Stack Access or Trunk Port

```
Dell(conf)#interface port-channel 10
Dell(conf-if-po-10)#switchport
Dell(conf-if-po-10)#vlt-peer-lag port-channel 10
Dell(conf-if-po-10)#vlan-stack access
Dell(conf-if-po-10)#no shutdown
```

```
Dell#show running-config interface port-channel 10
!
interface Port-channel 10
  no ip address
  switchport
  vlan-stack access
  vlt-peer-lag port-channel 10
  no shutdown
Dell#
```

```
Dell(conf)#interface port-channel 20
Dell(conf-if-po-20)#switchport
Dell(conf-if-po-20)#vlt-peer-lag port-channel 20
Dell(conf-if-po-20)#vlan-stack trunk
Dell(conf-if-po-20)#no shutdown
```

```
Dell#show running-config interface port-channel 20
!
interface Port-channel 20
  no ip address
  switchport
  vlan-stack trunk
```

```
vlt-peer-lag port-channel 20  
no shutdown  
Dell#
```

Configure the VLAN as VLAN-Stack VLAN and add the VLT LAG as members to the VLAN

```
Dell(conf)#interface vlan 50  
Dell(conf-if-vl-50)#vlan-stack compatible  
Dell(conf-if-vl-50-stack)#member port-channel 10  
Dell(conf-if-vl-50-stack)#member port-channel 20  
Dell(conf-if-vl-50-stack)#
```

```
Dell#show running-config interface vlan 50  
!  
interface Vlan 50  
  vlan-stack compatible  
  member Port-channel 10,20  
  shutdown  
Dell#
```

Verify that the Port Channels used in the VLT Domain are Assigned to the VLAN-Stack VLAN

Virtual Router Redundancy Protocol (VRRP)

Virtual router redundancy protocol (VRRP) is supported on the MXL switch platform.

VRRP Overview

VRRP is designed to eliminate a single point of failure in a statically routed network.

VRRP specifies a MASTER router that owns the next hop IP and MAC address for end stations on a local area network (LAN). The MASTER router is chosen from the virtual routers by an election process and forwards packets sent to the next hop IP address. If the MASTER router fails, VRRP begins the election process to choose a new MASTER router and that new MASTER continues routing traffic.

VRRP uses the virtual router identifier (VRID) to identify each virtual router configured. The IP address of the MASTER router is used as the next hop address for all end stations on the LAN. The other routers the IP addresses represent are BACKUP routers.

VRRP packets are transmitted with the virtual router MAC address as the source MAC address. The MAC address is in the following format: 00-00-5E-00-01-{VRID}. The first three octets are unchangeable. The next two octets (00-01) indicate the address block assigned to the VRRP protocol, and are unchangeable. The final octet changes depending on the VRRP virtual router identifier and allows for up to 255 VRRP routers on a network.

The following example shows a typical network configuration using VRRP. Instead of configuring the hosts on the network 10.10.10.0 with the IP address of either Router A or Router B as their default router; their default router is the IP address configured on the virtual router. When any host on the LAN segment wants to access the Internet, it sends packets to the IP address of the virtual router.

In the following example, Router A is configured as the MASTER router. It is configured with the IP address of the virtual router and sends any packets addressed to the virtual router through interface GigabitEthernet 1/1 to the Internet. As the BACKUP router, Router B is also configured with the IP address of the virtual router. If, for any reason, Router A becomes unavailable, VRRP elects a new MASTER Router. Router B assumes the duties of Router A and becomes the MASTER router. At that time, Router B responds to the packets sent to the virtual IP address.

All workstations continue to use the IP address of the virtual router to address packets destined to the Internet. Router B receives and forwards them on interface GigabitEthernet 10/1. Until Router A resumes operation, VRRP allows Router B to provide uninterrupted service to the users on the LAN segment accessing the Internet.

For more detailed information about VRRP, refer to *RFC 2338, Virtual Router Redundancy Protocol*.




Figure 135. Basic VRRP Configuration

VRRP Benefits

With VRRP configured on a network, end-station connectivity to the network is not subject to a single point-of-failure. End-station connections to the network are redundant and are not dependent on internal gateway protocol (IGP) protocols to converge or update routing tables.

VRRP Implementation

The MXL 10/40GbE Switch supports a total of 2000 VRRP groups on a switch and 255 VRRP groups per interface

Within a single VRRP group, up to 12 virtual IP addresses are supported. Virtual IP addresses can belong to the primary or secondary IP address' subnet configured on the interface. You can ping all the virtual IP addresses configured on the Master VRRP router from anywhere in the local subnet.

Default VRRP settings may affect the maximum number of groups that you can configure and work efficiently, as a result of hardware throttling VRRP advertisement packets reaching the CP on the MXL

switch. To avoid throttling VRRP advertisement packets, Dell Networking recommends increasing the VRRP advertisement interval to a value higher than the default value of 1 second.

△ CAUTION: Increasing the advertisement interval increases the VRRP Master dead interval, resulting in an increased failover time for Master/Backup election. Take caution when increasing the advertisement interval, as the increased dead interval may cause packets to be dropped during that switch-over time.

The recommendations in the following table vary, depending on several factors; for example, address resolution protocol (ARP) broadcasts, IP broadcasts, or spanning tree protocol (STP) before changing the advertisement interval. When the number of packets processed by the CP processor increases or decreases based on the dynamics of the network, the advertisement intervals in may increase or decrease accordingly.

Table 71. Recommended VRRP Advertise Intervals

Recommended Advertise Interval		Groups/Interface
Less than 250	1 second	255
Between 250 and 450	2–3 seconds	255
Between 450 and 600	3–4 seconds	255

VRRP Configuration

By default, VRRP is not configured.

Configuration Task List

The following list specifies the configuration tasks for VRRP.

- [Creating a Virtual Router](#) (mandatory)
- [Configuring the VRRP Version for an IPv4 Group](#) (optional)
- [Assign Virtual IP addresses](#) (mandatory)
- [Setting VRRP Group \(Virtual Router\) Priority](#) (optional)
- [Configuring VRRP Authentication](#) (optional)
- [Disabling Preempt](#) (optional)
- [Changing the Advertisement Interval](#) (optional)
- [Track an Interface or Object](#) (optional)
- [Setting VRRP Initialization Delay](#) (optional)

For a complete listing of all commands related to VRRP, refer to *Dell Networking OS Command Line Reference Guide*.

Creating a Virtual Router

To enable VRRP, create a virtual router. In Dell, the virtual router identifier (VRID) identifies a VRRP group. To enable or delete a virtual router, use the following commands.

- Create a virtual router for that interface with a VRID.

INTERFACE mode

```
vrrp-group vrid
```

The VRID range is from 1 to 255.



NOTE: The interface must already have a primary IP address defined and be enabled, as shown in the second example.

- Delete a VRRP group.

INTERFACE mode

```
no vrrp-group vrid
```

Example of Configuring VRRP

Example of Verifying the VRRP Configuration

```
Dell(conf)#int tengig 1/1
Dell(conf-if-te-1/1)#vrrp-group 111
Dell(conf-if-te-1/1-vrid-111)#
```

```
Dell(conf-if-te-1/1)#show conf
!
interface Tengigabitethernet 1/1
  ip address 10.10.10.1/24
!
  vrrp-group 111
  no shutdown
Dell(conf-if-te-1/1)#
```

Configuring the VRRP Version for an IPv4 Group

For IPv4, you can configure a VRRP group to use one of the following VRRP versions:

- VRRPv2 as defined in RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*
- VRRPv3 as defined in RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6*

You can also migrate a IPv4 group from VRRPv2 to VRRP3.

To configure the VRRP version for IPv4, use the **version** command in INTERFACE mode.

Example: Configuring VRRP to Use Version 3


The following example configures the IPv4 VRRP 100 group to use VRRP protocol version 3.

You can use the **version both** command in INTERFACE mode to migrate from VRRPv2 to VRRPv3. When you set the VRRP version to **both**, the switch sends only VRRPv3 advertisements but can receive VRRPv2 or VRRPv3 packets.


To migrate an IPv4 VRRP group from VRRPv2 to VRRPv3:

1. Set the switches with the lowest priority to “both”.

2. Set the switch with the highest priority to version to 3.
3. Set all the switches from **both** to version 3.

 **NOTE:** Do not run VRRP version 2 and version 3 in the same group for an extended period of time

Example: Migrating an IPv4 VRRP Group from VRRPv2 to VRRPv3

 **NOTE:** Carefully following this procedure, otherwise you might introduce dual master switches issues.

To migrate an IPv4 VRRP Group from VRRPv2 to VRRPv3:

1. Set the backup switches to VRRP version to both.
2. Set the master switch to VRRP protocol version 3.
3. Set the backup switches to version 3.

Assign Virtual IP addresses

Virtual routers contain virtual IP addresses configured for that VRRP group (VRID). A VRRP group does not transmit VRRP packets until you assign the Virtual IP address to the VRRP group.

To activate a VRRP group on an interface (so that VRRP group starts transmitting VRRP packets), configure at least one virtual IP address in a VRRP group. The virtual IP address is the IP address of the virtual router and does not require the IP address mask.

You can configure up to 12 virtual IP addresses on a single VRRP group (VRID).

The following rules apply to virtual IP addresses:

- The virtual IP addresses must be in the same subnet as the primary or secondary IP addresses configured on the interface. Though a single VRRP group can contain virtual IP addresses belonging to multiple IP subnets configured on the interface, Dell Networking recommends configuring virtual IP addresses belonging to the same IP subnet for any one VRRP group.
 - For example, an interface (on which you enable VRRP) contains a primary IP address of 50.1.1.1/24 and a secondary IP address of 60.1.1.1/24. The VRRP group (VRID 1) must contain virtual addresses belonging to either subnet 50.1.1.0/24 or subnet 60.1.1.0/24, but not from both subnets (though the system allows the same).
- If the virtual IP address and the interface's primary/secondary IP address are the same, the priority on that VRRP group **MUST** be set to 255. The interface then becomes the OWNER router of the VRRP group and the interface's physical MAC address is changed to that of the owner VRRP group's MAC address.
- If you configure multiple VRRP groups on an interface, only one of the VRRP Groups can contain the interface primary or secondary IP address.

Configuring a Virtual IP Address

To configure a virtual IP address, use the following commands.

1. Configure a VRRP group.
INTERFACE mode

`vrrp-group vrrp-id`

The VRID range is from 1 to 255.

2. Configure virtual IP addresses for this VRID.

INTERFACE -VRID mode

```
virtual-address ip-address1 [...ip-address12]
```


The range is up to 12 addresses.

Example of the `virtual-address` Command

Example of Verifying the Virtual IP Address Configuration

Example of Verifying the VRRP Group Priority

```
Dell(conf-if-te-1/1-vrid-111)#virtual-address 10.10.10.1
Dell(conf-if-te-1/1-vrid-111)#virtual-address 10.10.10.2
Dell(conf-if-te-1/1-vrid-111)#virtual-address 10.10.10.3
Dell(conf-if-te-1/1-vrid-111)#
```

 **NOTE:** In the following example, the primary IP address and the virtual IP addresses are on the same subnet.

```
Dell(conf-if-te-1/1)#show conf
!
interface Tengigabitethernet 1/1
  ip address 10.10.10.1/24
!
vrrp-group 111
  priority 255
  virtual-address 10.10.10.1
  virtual-address 10.10.10.2
  virtual-address 10.10.10.3
!
vrrp-group 222
no shutdown
Dell(conf-if-te-1/1)#
```

The following example shows the same VRRP group (VRID 111) configured on multiple interfaces on different subnets.

```
Dell#do show vrrp
-----
Tengigabitethernet 1/1, VRID: 111, Net: 10.10.10.1
State: Master, Priority: 255, Master: 10.10.10.1 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Bad pkts rcvd: 0, Adv sent: 1768, Gratuitous ARP sent: 5
Virtual MAC address:
  00:00:5e:00:01:6f
Virtual IP address:
  10.10.10.1 10.10.10.2 10.10.10.3 10.10.10.10
Authentication: (none)
-----
Tengigabitethernet 1/2, VRID: 111, Net: 10.10.2.1
State: Master, Priority: 100, Master: 10.10.2.1 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Bad pkts rcvd: 0, Adv sent: 27, Gratuitous ARP sent: 2
Virtual MAC address:
  00:00:5e:00:01:6f
Virtual IP address:
  10.10.2.2 10.10.2.3
```

Authentication: (none)
Dell#

When the VRRP process completes its initialization, the State field contains either Master or Backup.

Setting VRRP Group (Virtual Router) Priority

Setting a virtual router priority to 255 ensures that router is the “owner” virtual router for the VRRP group. VRRP elects the MASTER router by choosing the router with the highest priority.

The default priority for a virtual router is **100**. The higher the number, the higher the priority. If the MASTER router fails, VRRP begins the election process to choose a new MASTER router based on the next-highest priority.

If two routers in a VRRP group come up at the same time and have the same priority value, the interface’s physical IP addresses are used as tie-breakers to decide which is MASTER. The router with the higher IP address becomes MASTER.

To configure the VRRP group’s priority, use the following command.

- Configure the priority for the VRRP group.

```
INTERFACE -VRID mode
```

```
priority priority
```

The range is from 1 to 255.

The default is **100**.

Example of the `priority` Command

Example of Verifying the VRRP Group Priority


```
Dell(conf-if-te-1/2)#vrrp-group 111  
Dell(conf-if-te-1/2-vrid-111)#priority 125
```

```
Dell#show vrrp
```

```
-----  
Tengigabitethernet 1/1, VRID: 111, Net: 10.10.10.1  
State: Master, Priority: 255, Master: 10.10.10.1 (local)  
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec  
Adv rcvd: 0, Bad pkts rcvd: 0, Adv sent: 2343, Gratuitous ARP sent: 5  
Virtual MAC address:  
  00:00:5e:00:01:6f  
Virtual IP address:  
  10.10.10.1 10.10.10.2 10.10.10.3 10.10.10.10  
Authentication: (none)  
-----  
Tengigabitethernet 1/2, VRID: 111, Net: 10.10.2.1  
State: Master, Priority: 125, Master: 10.10.2.1 (local)  
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec  
Adv rcvd: 0, Bad pkts rcvd: 0, Adv sent: 601, Gratuitous ARP sent: 2  
Virtual MAC address:  
  00:00:5e:00:01:6f  
Virtual IP address:  
  10.10.2.2 10.10.2.3  
Authentication: (none)  
Dell(conf)#
```

Configuring VRRP Authentication

Simple authentication of VRRP packets ensures that only trusted routers participate in VRRP processes. When you enable authentication, the Dell Networking OS includes the password in its VRRP transmission. The receiving router uses that password to verify the transmission.

 **NOTE:** You must configure all virtual routers in the VRRP group the same: you must enable authentication with the same password or authentication is disabled.

To configure simple authentication, use the following command.

- Configure a simple text password.

INTERFACE-VRID mode

```
authentication-type simple [encryption-type] password
```

Parameters:

- *encryption-type*: 0 indicates unencrypted; 7 indicates encrypted.
- *password*: plain text.

Example of authentication-type Command

Example of Verifying the Configuration of VRRP Authentication

The bold section shows the encryption type (encrypted) and the password.

```
Dell(conf-if-te-1/1-vrid-111)#authentication-type ?
Dell(conf-if-te-1/1-vrid-111)#authentication-type simple 7 force10
```


The bold section shows the encrypted password.

```
Dell(conf-if-te-1/1-vrid-111)#show conf
!
vrrp-group 111
 authentication-type simple 7 387a7f2df5969da4
 priority 255
 virtual-address 10.10.10.1
 virtual-address 10.10.10.2
 virtual-address 10.10.10.3
 virtual-address 10.10.10.10
Dell(conf-if-te-1/1-vrid-111)#
```

Disabling Preempt

The `preempt` command is enabled by default. The command forces the system to change the MASTER router if another router with a higher priority comes online.

Prevent the BACKUP router with the higher priority from becoming the MASTER router by disabling preempt.

 **NOTE:** You must configure all virtual routers in the VRRP group the same: you must configure all with preempt enabled or configure all with preempt disabled.

Because preempt is enabled by default, disable the preempt function with the following command.

- Prevent any BACKUP router with a higher priority from becoming the MASTER router.

INTERFACE-VRID mode

```
no preempt
```

Example of Disabling Preempt

Example of Verifying Preempt is Disabled

Re-enable preempt by entering the `preempt` command. When you enable preempt, it does not display in the `show` commands, because it is a default setting.


```
Dell(conf-if-te-1/1)#vrrp-group 111
Dell(conf-if-te-1/1-vrid-111)#no preempt
Dell(conf-if-te-1/1-vrid-111)#show conf

Dell(conf-if-te-1/1-vrid-111)#show conf
!
vrrp-group 111
 authentication-type simple 7 387a7f2df5969da4
  no preempt
  priority 255
  virtual-address 10.10.10.1
  virtual-address 10.10.10.2
  virtual-address 10.10.10.3
  virtual-address 10.10.10.10
Dell(conf-if-te-1/1-vrid-111)#
```

Changing the Advertisement Interval

By default, the MASTER router transmits a VRRP advertisement to all members of the VRRP group every one second, indicating it is operational and is the MASTER router.

If the VRRP group misses three consecutive advertisements, the election process begins and the BACKUP virtual router with the highest priority transitions to MASTER.

 **NOTE:** To avoid throttling VRRP advertisement packets, Dell Networking recommends increasing the VRRP advertisement interval to a value higher than the default value of one second. If you do change the time interval between VRRP advertisements on one router, change it on all participating routers.

If are using VRRP version 2, you must configure the timer values in multiple of whole seconds. For example a timer value of 3 seconds or 300 centiseocs are valid and equivalent. However, a time value of 50 centiseocs is invalid because it not a multiple of 1 second. If you are using VRRP version 3, you must configure the timer values in multiples of 25 centiseocs.

If you are configured for VRRP version 2, the timer values must be in multiples of whole seconds. For example, timer value of 3 seconds or 300 centiseocs are valid and equivalent. However, a timer value of 50 centiseocs is invalid because it not is not multiple of 1 second.

If are using VRRP version 3, you must configure the timer values in multiples of 25 centiseocs.

To change the advertisement interval in seconds or centiseocs, use the following command. A centiseocs is 1/100 of a second.

1. Change the advertisement interval setting. INTERFACE-VRID mode `advertise-interval seconds` The range is from 1 to 255 seconds. The default is **1 second**.
2. For VRRPv3, change the advertisement centiseocs interval setting. INTERFACE-VRID mode `advertise-interval centiseocs centiseocs`

The range is from 25 to 4075 centiseocs in units of 25 centiseocs.

The default is **100 centiseocs**.

Example of the `advertise-interval` Command

Example of Verifying the Configured Advertisement Interval

The following example shows how to change the advertise interval using the `advertise-interval` command.

```
Dell(conf-if-te-1/1)#vrrp-group 111
Dell(conf-if-te-1/1-vrid-111)#advertise-interval 10
Dell(conf-if-te-1/1-vrid-111)#
```

The following example shows how to verify the advertise interval change using the `show conf` command

```
Dell(conf-if-te-1/1-vrid-111)#show conf
!
vrrp-group 111
  advertise-interval 10
  authentication-type simple 7 387a7f2df5969da4
  no preempt
  priority 255
  virtual-address 10.10.10.1
  virtual-address 10.10.10.2
  virtual-address 10.10.10.3
  virtual-address 10.10.10.10
Dell(conf-if-te-1/1-vrid-111)#
```

Track an Interface or Object

You can set the Dell Networking OS to monitor the state of any interface according to the virtual group.


Each VRRP group can track up to 12 interfaces and up to 20 additional objects, which may affect the priority of the VRRP group. If the tracked interface goes down, the VRRP group's priority decreases by a default value of **10** (also known as *cost*). If the tracked interface's state goes up, the VRRP group's priority increases by 10.

The lowered priority of the VRRP group may trigger an election. As the Master/Backup VRRP routers are selected based on the VRRP group's priority, tracking features ensure that the best VRRP router is the Master for that group. The sum of all the costs of all the tracked interfaces must be less than the configured priority on the VRRP group. If the VRRP group is configured as Owner router (priority 255), tracking for that group is disabled, irrespective of the state of the tracked interfaces. The priority of the owner group always remains at 255.

For a virtual group, you can track the line-protocol state or the routing status of any of the following interfaces with the `interface interface` parameter:


- 40 Gigabit Ethernet: enter `fortygigabitethernet slot/port` in the track interface command (shown in the following example).
- 10 Gigabit Ethernet: enter `tengigabitethernet slot/port`.
- Port channel: enter `port-channel number`.
- VLAN: enter `vlan vlan-id` where valid VLAN IDs are from 1 to 4094.

For a virtual group, you can also track the status of a configured object (the `track object-id` command) by entering its object number.

-  **NOTE:** You can configure a tracked object for a VRRP group (using the `track object-id` command in INTERFACE-VRID mode) before you actually create the tracked object (using a `track object-id` command in CONFIGURATION mode). However, no changes in the VRRP group's priority occur until the tracked object is defined and determined to be down.

Tracking an Interface

To track an interface, use the following commands.

-  **NOTE:** The sum of all the costs for all tracked interfaces must be less than the configured priority of the VRRP group.

- Monitor an interface and, optionally, set a value to be subtracted from the interface's VRRP group priority.

INTERFACE-VRID mode

```
track interface [priority-cost cost]
```

The cost range is from 1 to 254.

The default is **10**.

- (Optional) Display the configuration and the UP or DOWN state of tracked objects, including the client (VRRP group) that is tracking an object's state.

EXEC mode or EXEC Privilege mode

```
show track
```

- (Optional) Display the configuration and the UP or DOWN state of tracked interfaces and objects in VRRP groups, including the time since the last change in an object's state.

EXEC mode or EXEC Privilege mode

```
show vrrp
```

- (Optional) Display the configuration of tracked objects in VRRP groups on a specified interface.

EXEC mode or EXEC Privilege mode

```
show running-config interface interface
```

Example of the `track` Command

Example of Verifying the Tracking Configuration

Example of Viewing Tracking Status

Example of Viewing VRRP Configuration on an Interface

```
Dell(conf-if-te-1/1)#vrrp-group 111
Dell(conf-if-te-1/1-vrid-111)#track tengigabitethernet 1/2
Dell(conf-if-te-1/1-vrid-111)#
```

```
Dell(conf-if-te-1/1-vrid-111)#show conf
!
vrrp-group 111
  advertise-interval 10
  authentication-type simple 7 387a7f2df5969da4
  no preempt
  priority 255
  track Tengigabitethernet 1/2
  virtual-address 10.10.10.1
```

```
virtual-address 10.10.10.2
virtual-address 10.10.10.3
virtual-address 10.10.10.10
Dell(conf-if-te-1/1-vrid-111)#
```

```
Dell#show track
```

```
Track 2
  IPv6 route 2040::/64
  metric threshold Metric threshold is Up (STATIC/0/0)
    5 changes, last change 00:02:16
  Metric threshold down 255 up 254
  First-hop interface is GigabitEthernet 13/2
Tracked by:
VRRP GigabitEthernet 7/30 IPv6 VRID 1
```

```
Track 3
  IPv6 route 2050::/64 reachability
  Reachability is Up (STATIC)
    5 changes, last change 00:02:16
  First-hop interface is GigabitEthernet 13/2
Tracked by:
VRRP GigabitEthernet 7/30 IPv6 VRID 1
```


```
Dell#show running-config interface tengigabitethernet 1/3
!
interface TenGigabitEthernet 1/3
ip address 10.1.1.1/24
!
vrrp-group 21
virtual-address 10.1.1.2
no shutdown
Dell#
```

Setting VRRP Initialization Delay

When configured, VRRP is enabled immediately upon system reload or boot.

You can delay VRRP initialization to allow the IGP and EGP protocols to be enabled prior to selecting the VRRP Master. This delay ensures that VRRP initializes with no errors or conflicts. You can configure the delay for up to 15 minutes, after which VRRP enables normally.

Set the delay timer on individual interfaces. The delay timer is supported on all physical interfaces, VLANs, and LAGs.

 **NOTE:** When you reload a node that contains VRRP configuration and is enabled for VLT, Dell Networking recommends that you configure the reload timer by using the `vrrp delay reload` command to ensure that VRRP is functional. Otherwise, when you reload a VLT node configured for VRRP, the local destination address is not seen on the reloaded node causing suboptimal routing.

When you configure both CLIs, the later timer rules VRRP enabling. For example, if you set `vrrp delay reload 600` and `vrrp delay minimum 300`, the following behavior occurs:

- When the system reloads, VRRP waits 600 seconds (10 minutes) to bring up VRRP on all interfaces that are up and configured for VRRP.
- When an interface comes up and becomes operational, the system waits 300 seconds (5 minutes) to bring up VRRP on that interface.

To set the delay time for VRRP initialization, use the following commands.

- Set the delay time for VRRP initialization on an individual interface.

INTERFACE mode

```
vrrp delay minimum seconds
```

This time is the gap between an interface coming up and being operational, and VRRP enabling.

The seconds range is from 0 to 900.

The default is **0**.

- Set the delay time for VRRP initialization on all the interfaces in the system configured for VRRP.

INTERFACE mode

```
vrrp delay reload seconds
```

This time is the gap between system boot up completion and VRRP enabling.

The seconds range is from 0 to 900.

The default is **0**.

Sample Configurations

Before you set up VRRP, review the following sample configurations.

VRRP for an IPv4 Configuration

The following configuration shows how to enable IPv4 VRRP. This example does not contain comprehensive directions and is intended to provide guidance for only a typical VRRP configuration. You can copy and paste from the example to your CLI. To support your own IP addresses, interfaces, names, and so on, be sure that you make the necessary changes. The VRRP topology was created using the CLI configuration shown in the following example.


```

R2#show vrrp
-----
TenGigabitEthernet 2/31, VRID: 99, Net: 10.1.1.1
State: Master, Priority: 100, Master: 10.1.1.1 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Bad pkts rcvd: 0, Adv sent: 661, Gratuitous ARP sent: 1
Virtual MAC address:
00:00:5e:00:01:63
Virtual IP address:
10.1.1.3
Authentication: (none)
R2#

R3#show vrrp
-----
TenGigabitEthernet 3/21, VRID: 99, Net: 10.1.1.1
State: Backup, Priority: 100, Master: 10.1.1.1 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 331, Bad pkts rcvd: 0, Adv sent: 0, Gratuitous ARP sent: 0
Virtual MAC address:
00:00:5e:00:01:63
Virtual IP address:
10.1.1.3
Authentication: (none)
R3#

```

State Master: R2 was the first interface configured with VRRP

Virtual MAC is automatically assigned and is the same on both Routers

State Backup: R3 was the second interface configured with VRRP

Virtual MAC is automatically assigned and is the same on both Routers




Figure 136. VRRP for IPv4 Topology

Example of Configuring VRRP for IPv4

```

R2(conf)#int tengig 2/31
R2(conf-if-te-2/31)#ip address 10.1.1.1/24
R2(conf-if-te-2/31)#vrrp-group 99
R2(conf-if-te-2/31-vrid-99)#priority 200
R2(conf-if-te-2/31-vrid-99)#virtual 10.1.1.3
R2(conf-if-te-2/31-vrid-99)#no shut
R2(conf-if-te-2/31)#show conf
!
interface Tengigabitethernet 2/31
 ip address 10.1.1.1/24
!
 vrrp-group 99
 priority 200
 virtual-address 10.1.1.3
 no shutdown

```


```

R2(conf-if-te-2/31)#end
R2#show vrrp
-----
Tengigabitethernet 2/31, VRID: 99, Net: 10.1.1.1
State: Master, Priority: 200, Master: 10.1.1.1 (local)
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 0, Bad pkts rcvd: 0, Adv sent: 817, Gratuitous ARP sent: 1
Virtual MAC address:
  00:00:5e:00:01:63
Virtual IP address:
  10.1.1.3
Authentication: (none)
R2#
Router 3
R3(conf)#int tengig 3/21
R3(conf-if-te-3/21)#ip address 10.1.1.2/24
R3(conf-if-te-3/21)#vrrp-group 99
R3(conf-if-te-3/21-vrid-99)#virtual 10.1.1.3
R3(conf-if-te-3/21-vrid-99)#no shut
R3(conf-if-te-3/21)#show conf
!
interface Tengigabitethernet 3/21
ip address 10.1.1.2/24
!
vrrp-group 99
  virtual-address 10.1.1.3
  no shutdown
R3(conf-if-te-3/21)#end
R3#show vrrp
-----
Tengigabitethernet 3/21, VRID: 99, Net: 10.1.1.2
State: Backup, Priority: 100, Master: 10.1.1.1
Hold Down: 0 sec, Preempt: TRUE, AdvInt: 1 sec
Adv rcvd: 698, Bad pkts rcvd: 0, Adv sent: 0, Gratuitous ARP sent: 0
Virtual MAC address:
  00:00:5e:00:01:63
Virtual IP address:
  10.1.1.3
Authentication: (none)

```

Standards Compliance

This chapter describes standards compliance for Dell Networking products.

 **NOTE:** Unless noted, when a standard cited here is listed as supported by the Dell Networking Operating System (OS), the system also supports predecessor standards. One way to search for predecessor standards is to use the <http://tools.ietf.org/> website. Click "Browse and search IETF documents," enter an RFC number, and inspect the top of the resulting document for obsolescence citations to related RFCs.

IEEE Compliance

The following is a list of IEEE compliance.

802.1AB	LLDP
802.1D	Bridging, STP
802.1p	L2 Prioritization
802.1Q	VLAN Tagging, Double VLAN Tagging, GVRP
802.1s	MSTP
802.1w	RSTP
802.3ac	Frame Extensions for VLAN Tagging
802.3ad	Link Aggregation with LACP
802.3ae	10 Gigabit Ethernet (10GBASE-W, 10GBASE-X)
802.3ak	10 Gigabit Ethernet (10GBASE-CX4)
802.3i	Ethernet (10BASE-T)
802.3u	Fast Ethernet (100BASE-FX, 100BASE-TX)
802.3x	Flow Control
802.1Qaz	Enhanced Transmission Selection
802.1Qbb	Priority-based Flow Control
ANSI/TIA-1057	LLDP-MED
Dell Networking	FRRP (Force10 Redundant Ring Protocol)
802.1w	PVST+

SFF-8431	SFP+ Direct Attach Cable (10GSFP+Cu)
MTU	12,000 bytes

RFC and I-D Compliance

The Dell Networking OS supports the following standards. The standards are grouped by related protocol. The columns showing support by platform indicate which version of Dell Networking OS first supports the standard.

General Internet Protocols

The following table lists the Dell Networking OS support per platform for general internet protocols.

Table 72. General Internet Protocols

RFC#	Full Name
768	User Datagram Protocol
793	Transmission Control Protocol
854	Telnet Protocol Specification
959	File Transfer Protocol (FTP)
1321	The MD5 Message-Digest Algorithm
1350	The TFTP Protocol (Revision 2)
1661	The Point-to-Point Protocol (PPP)
1989	PPP Link Quality Monitoring
1990	The PPP Multilink Protocol (MP)
1994	PPP Challenge Handshake Authentication Protocol (CHAP)
2474	Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers
2698	A Two Rate Three Color Marker
3164	The BSD syslog Protocol
draft-ietf-bfd -base-03	Bidirectional Forwarding Detection

General IPv4 Protocols

The following table lists the Dell Networking OS support per platform for general IPv4 protocols.

Table 73. General IPv4 Protocols

RFC#	Full Name
791	Internet Protocol
792	Internet Control Message Protocol
826	An Ethernet Address Resolution Protocol
1027	Using ARP to Implement Transparent Subnet Gateways
1035	DOMAIN NAMES - IMPLEMENTATION AND SPECIFICATION (client)
1042	A Standard for the Transmission of IP Datagrams over IEEE 802 Networks
1191	Path MTU Discovery
1305	Network Time Protocol (Version 3) Specification, Implementation and Analysis
1519	Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy
1542	Clarifications and Extensions for the Bootstrap Protocol
1812	Requirements for IP Version 4 Routers
2131	Dynamic Host Configuration Protocol
2338	Virtual Router Redundancy Protocol (VRRP)
3021	Using 31-Bit Prefixes on IPv4 Point-to-Point Links
3046	DHCP Relay Agent Information Option
3069	VLAN Aggregation for Efficient IP Address Allocation
3128	Protection Against a Variant of the Tiny Fragment Attack

Border Gateway Protocol (BGP)

The following table lists the Dell Networking OS support per platform for BGP protocols.

Table 74. Border Gateway Protocol (BGP)

RFC#	Full Name
1997	BGP ComAmturnbituitees
2385	Protection of BGP Sessions via the TCP MD5 Signature Option
2439	BGP Route Flap Damping
2796	BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)
2842	Capabilities Advertisement with BGP-4
2858	Multiprotocol Extensions for BGP-4
2918	Route Refresh Capability for BGP-4
3065	Autonomous System Confederations for BGP
4360	BGP Extended Communities Attribute
4893	BGP Support for Four-octet AS Number Space
5396	Textual Representation of Autonomous System (AS) Numbers
draft-ietf-idrbgp4- 20	A Border Gateway Protocol 4 (BGP-4)
draft-ietf-idrrestart- 06	Graceful Restart Mechanism for BGP

Open Shortest Path First (OSPF)

The following table lists the Dell Networking OS support per platform for OSPF protocol.

Table 75. Open Shortest Path First (OSPF)

RFC#	Full Name
1587	The OSPF Not-So-Stubby Area (NSSA) Option
2154	OSPF with Digital Signatures
2328	OSPF Version 2
2370	The OSPF Opaque LSA Option
3623	Graceful OSPF Restart
4222	Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance

Routing Information Protocol (RIP)

The following table lists the Dell Networking OS support per platform for RIP protocol.

Table 76. Routing Information Protocol (RIP)

RFC#	Full Name
1058	Routing Information Protocol
2453	RIP Version 2

Network Management

The following table lists the Dell Networking OS support per platform for network management protocol.

Table 77. Network Management

RFC#	Full Name
1155	Structure and Identification of Management Information for TCP/IP-based Internets
1156	Management Information Base for Network Management of TCP/IP-based internets
1157	A Simple Network Management Protocol (SNMP)
1212	Concise MIB Definitions
1215	A Convention for Defining Traps for use with the SNMP
1493	Definitions of Managed Objects for Bridges [except for the dot1dTpLearnedEntryDiscards object]
1724	RIP Version 2 MIB Extension
1850	OSPF Version 2 Management Information Base
1901	Introduction to Community-based SNMPv2
2011	SNMPv2 Management Information Base for the Internet Protocol using SMIv2
2012	SNMPv2 Management Information Base for the Transmission Control Protocol using SMIv2
2013	SNMPv2 Management Information Base for the User Datagram Protocol using SMIv2
2024	Definitions of Managed Objects for Data Link Switching using SMIv2
2096	IP Forwarding Table MIB

RFC#	Full Name
2570	Introduction and Applicability Statements for Internet Standard Management Framework
2571	An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks
2572	Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
2574	User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)
2575	View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)
2576	Coexistence Between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework
2578	Structure of Management Information Version 2 (SMIv2)
2579	Textual Conventions for SMIv2
2580	Conformance Statements for SMIv2
2618	RADIUS Authentication Client MIB, except the following four counters: radiusAuthClientInvalidServerAddresses radiusAuthClientMalformedAccessResponses radiusAuthClientUnknownTypes radiusAuthClientPacketsDropped
2698	A Two Rate Three Color Marker
3635	Definitions of Managed Objects for the Ethernet-like Interface Types
2674	Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions
2787	Definitions of Managed Objects for the Virtual Router Redundancy Protocol
2819	Remote Network Monitoring Management Information Base: Ethernet Statistics Table, Ethernet History Control Table, Ethernet History Table, Alarm Table, Event Table, Log Table
2863	The Interfaces Group MIB
2865	Remote Authentication Dial In User Service (RADIUS)
3273	Remote Network Monitoring Management Information Base for High Capacity Networks (64 bits): Ethernet

RFC#	Full Name
	Statistics High-Capacity Table, Ethernet History High-Capacity Table
3416	Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)
3418	Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)
3434	Remote Monitoring MIB Extensions for High Capacity Alarms, High-Capacity Alarm Table (64 bits)
3580	IEEE 802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines
3815	Definitions of Managed Objects for the Multiprotocol Label Switching (MPLS), Label Distribution Protocol (LDP)
4001	Textual Conventions for Internet Network Addresses
4292	IP Forwarding Table MIB
4750	OSPF Version 2 Management Information Base
4520	RMON v2 MIB
5060	Protocol Independent Multicast MIB
ANSI/TIA-1057	The LLDP Management Information Base extension module for TIA-TR41.4 Media Endpoint Discovery information
draft-grant-tacacs -02	The TACACS+ Protocol
draft-ietf-idr-bgp4 -mib-06	Definitions of Managed Objects for the Fourth Version of the Border Gateway Protocol (BGP-4) using SMIPv2
IEEE 802.1AB	Management Information Base module for LLDP configuration, statistics, local system data and remote systems data components.
IEEE 802.1AB	The LLDP Management Information Base extension module for IEEE 802.1 organizationally defined discovery information. (LLDP DOT1 MIB and LLDP DOT3 MIB)
IEEE 802.1AB	The LLDP Management Information Base extension module for IEEE 802.3 organizationally defined discovery information. (LLDP DOT1 MIB and LLDP DOT3 MIB)
ruzin-mstp-mib-0 2 (Traps)	Definitions of Managed Objects for Bridges with Multiple Spanning Tree Protocol
sFlow.org	sFlow Version 5
sFlow.org	sFlow Version 5 MIB
FORCE10-BGP4-V2-MIB	Force10 BGP MIB (draft-ietf-idr-bgp4-mibv2-05)

RFC#	Full Name
FORCE10-IF-EXTENSION-MIB	Force10 Enterprise IF Extension MIB (extends the Interfaces portion of the MIB-2 (RFC 1213) by providing proprietary SNMP OIDs for other counters displayed in the "show interfaces" output)
FORCE10-LINKAGG-MIB	Force10 Enterprise Link Aggregation MIB
FORCE10-COPY-CONFIG-MIB	Force10 File Copy MIB (supporting SNMP SET operation)
FORCE10-MONMIB	Force10 Monitoring MIB
FORCE10-PRODUCTS-MIB	Force10 Product Object Identifier MIB
FORCE10-SS-CHASSIS-MIB	Force10 S-Series Enterprise Chassis MIB
FORCE10-SMI	Force10 Structure of Management Information
FORCE10-SYSTEM-COMPONENT-MIB	Force10 System Component MIB (enables the user to view CAM usage information)
FORCE10-TC-MIB	Force10 Textual Convention
FORCE10-TRAP-ALARM-MIB	Force10 Trap Alarm MIB
FORCE10-FIPS NOOPING-MI B	Force10 FIP Snooping MIB (Based on T11-FCoE-MIB mentioned in FC-BB-5)
FORCE10-DCB -MIB	Force10 DCB MIB
IEEE 802.1Qaz	Management Information Base extension module for IEEE 802.1 organizationally defined discovery information (LDP-EXT-DOT1-DCBX-MIB)
IEEE 802.1Qbb	Priority-based Flow Control module for managing IEEE 802.1Qbb

MIB Location

You can find Force10 MIBs under the Force10 MIBs subhead on the Documentation page of iSupport:

<https://www.force10networks.com/csportal20/KnowledgeBase/Documentation.aspx>

You also can obtain a list of selected MIBs and their OIDs at the following URL:

https://www.force10networks.com/csportal20/MIBs/MIB_OIDs.aspx

Some pages of iSupport require a login. To request an iSupport account, go to:

<https://www.force10networks.com/CSPortal20/Support/AccountRequest.aspx>

If you have forgotten or lost your account information, contact Dell TAC for assistance.

FC Flex IO Modules

This part provides a generic, broad-level description of the operations, capabilities, and configuration commands of the Fiber Channel (FC) Flex IO module.

FC Flex IO Modules

This part provides a generic, broad-level description of the operations, capabilities, and configuration commands of the Fiber Channel (FC) Flex IO module.

Understanding and Working of the FC Flex IO Modules

This chapter provides a generic, broad-level description of the operations and functionality of the Fiber Channel (FC) Flex IO module, and contains the following sections:

- [FC Flex IO Modules Overview](#)
- [FC Flex IO Module Capabilities and Operations](#)
- [Guidelines for Working with FC Flex IO Modules](#)
- [Processing of Data Traffic](#)
- [Installing and Configuring the Switch](#)
- [Interconnectivity of FC Flex IO Modules with Cisco MDS Switches](#)

FC Flex IO Modules Overview

The Fibre Channel (FC) Flex IO module is supported on MXL 10/40GbE Switch and M I/O Aggregator (IOA). The MXL or IOA switch installed with the FC Flex IO module functions as a top-of-rack edge switch that supports converged enhanced Ethernet (CEE) traffic — Fibre Channel over Ethernet (FCoE) for storage, Interprocess Communication (IPC) for servers, and Ethernet local area network (LAN) (IP cloud) for data — as well as FC links to one or more storage area network (SAN) fabrics. Although the MXL 10/40GbE Switch and the I/O Aggregator can act as a FIP snooping bridge (FSB) to provide FCoE transit switch capabilities, the salient and significant advantage of deploying the FC Flex IO module is to enable more streamlined and cohesive FCoE N-port identifier virtualization (NPIV) proxy gateway functionalities. The NPIV proxy gateway (NPG) provides FCoE-FC bridging behavior.

The FC Flex IO module offers a rich, comprehensive set of FCoE functionalities on the M1000e chassis by splitting the Ethernet and Fibre Channel (FC) traffic at the edge of the chassis. The FC switches that are connected directly to the FC Flex IO module provide Fibre Channel capabilities because the FC Flex IO module does not support full fabric functionalities. With the separation of Ethernet and FC packets performed at the edge of the chassis itself, you can use the MXL 10/40GbE Switch or the Aggregator that contains an FC Flex IO module to connect to a SAN environment without the need for a separate ToR

switch to operate as NPIV proxy gateways. The MXL 10/40GbE Switch or the I/O Aggregator can function in NPIV proxy gateway mode when an FC Flex IO module is present or in the FIP snooping bridge (FSB) mode when all the ports are Ethernet ports.


The FC Flex IO module uses the same baseboard hardware of the MXL 10/40GbE Switch or the Aggregator and the M1000 chassis. You can insert the FC Flex IO module into any of the optional module slots of the MXL 10/40GbE Switch and it provides four FC ports per module. If you insert only one FC Flex IO module, four ports are supported; if you insert two FC Flex IO modules, eight ports are supported.

By installing an FC Flex IO module, you can enable the MXL 10/40GbE Switch and I/O Aggregator to directly connect to an existing FC SAN network. The FC Flex IO module uses the existing slots on the MXL 10/40GbE Switch and I/O Aggregator and provides four or eight FC ports up to speed of 8 GbE per second. You can connect all of the FC ports to the same FC SAN fabric to yield FC bandwidth of up to 64GB. It is possible to connect some of the ports to a different FC SAN fabric to provide access to multiple fabric devices.

In a typical Fibre Channel storage network topology, separate network interface cards (NICs) and host bus adapters (HBAs) on each server (two each for redundancy purposes) are connected to LAN and SAN networks respectively. These deployments typically include a ToR SAN switch in addition to a ToR LAN switch. By employing converged network adapters (CNAs) that the FC Flex IO module supports, CNAs are used to transmit FCoE traffic from the server instead of separate NIC and HBA devices. In such a scenario, you can determine whether the FC or SAN packets and the Ethernet or LAN packets must be split within the chassis or by using a ToR switch to perform this splitting.

If you want to segregate the LAN and SAN traffic within the chassis, you can employ switches such as the Dell M8428-k Converged 10GbE Switch or FC-only switches such as the Dell M5424 switch module. You can also use the S5000 Switch as a ToR switch to separate the LAN and SAN traffic at the ToR. By using the FC Flex IO module, you can optimally and effectively split the LAN and SAN traffic at the edge of the blade chassis itself. You can deploy the FC Flex IO module can be deployed in the enterprise and data center switching networks to leverage and derive the advantages of a converged Ethernet network.

The FC Flex IO module is not an FCF switch, but it offers FCoE capabilities from the server to the MXL or I/O Aggregator switches, and native FC capability in the uplink direction to the SAN switches. Although the FC Flex IO module does not support all of the FCF characteristics, such as full-blown name services or zone parameters, it presents the most flexible solution in interoperating with third-party switches that enable the splitting of LAN and SAN traffic. With the MXL 10/40GbE Switch and I/O Aggregator being well-established systems in the switch domain, you can install the FC Flex IO module to enhance and increase the converged Ethernet network performance and behavior. With the FC Flex IO module, the MXL 10/40GbE Switch and I/O Aggregator provide thirty-two 1GbE or 10 GbE server-facing ports and the option to add two FC Flex IO modules that offer up to 8 8Gb Fibre Channel ports for uplink traffic in addition to the fixed two 40GbE ports on the MXL 10/40GbE Switch and I/O Aggregator.

 **NOTE:** When an FC Flex IO module is inserted into an I/O Aggregator and the FC ports are in the operationally up state, you can configure the port speed of these FC ports as 2 Gbps, 4 Gbps, or 8 Gbps. In the chassis management controller (CMC) GUI, the FC port link speed is always shown as 10 Gbps, regardless of whether the port speed configured is 2 Gbps, 4 Gbps, or 8 Gbps.

You can configure one of the following upstream (fabric-facing) FC ports:

- Two 40GbE and eight 8GB FC ports
- Four 40GbE and four 8GB FC ports
- Two 40GbE, four 10GbE, and four 8GB FC ports

- Two 40GbE, four 10GBASE-T, and four 8GB FC ports

FC Flex IO Module Capabilities and Operations

The FC Flex IO module has the following characteristics:

- You can install one or two FC Flex IO modules on the MXL 10/40GbE Switch or I/O Aggregator. Each module supports four FC ports.
- Each port can operate in 2Gbps, 4Gbps, or 8Gbps of Fibre Channel speed.
- All ports on an FC Flex IO module can function in the NPIV mode that enables connectivity to FC switches or directors, and also to multiple SAN topologies.
- It automatically senses the current speed when the port link is up. Valid link speeds are 2 Gbps, 4 Gbps, and 8 Gbps.
- By default, the FC ports are configured to operate in N-port mode to connect to an F port on an FC switch in a fabric. You can apply only one FCoE map on an FC port. An N-Port is a port on the node of an FC device and is called a node port.
- There should be a maximum of 64 server fabric login (FLOGI) requests or fabric discovery (FDISC) requests per server MAC address before being forwarded by the FC Flex IO module to the FC core switch. Without user configuration, only 32 server login sessions are permitted for each server MAC address. To increase the total number of sessions to 64, use the `max_sessions` command.
- A distance of up to 300 meters is supported at 8 Gbps for Fibre Channel traffic.
- Multiple domains are supported in an NPIV proxy gateway (NPG).
- You cannot configure the MXL or Aggregator switches in Stacking mode if the switches contain the FC Flex IO module. Similarly, FC Flex IO modules do not function when you insert them in to a stack of MXL or Aggregator switches.
- If the switch does not contain FC Flex modules, you cannot create a stack, and a log message states that stacking is not supported unless the switches contain only FC Flex modules.

Guidelines for Working with FC Flex IO Modules

The following guidelines apply to the FC Flex IO module:

- All the ports of FC Flex IO modules operate in FC mode, and do not support Ethernet mode.
- FC Flex IO modules are not supported in the chassis management controller (CMC) GUI.
- The only supported FCoE functionality is NPIV proxy gateway. Configure the other FCoE services, such as name server, zone server, and login server on an external FC switch.
- With the FC Flex IO module, the MXL 10/40GbE Switch continues to support bare metal provisioning (BMP) on any Ethernet port. BMP is not supported on FC ports. BMP improves accessibility to the MXL 10/40GbE Switch by automatically loading pre-defined configurations and boot images that are stored in file servers. You can use BMP on a single switch or on multiple switches.
- FC Flex IOM module is a field-replaceable unit (FRU). Its memory type is electrically erasable programmable read-only memory (EEPROM), which enables it to save manufacturing information, such as the serial number. It is hot-swappable, assuming that the module that is removed is replaced by the same type of module in that same slot.

- The FC Flex IO does not have persistent storage for any runtime configuration. All the persistent storage for runtime configuration is on the MXL and IOA baseboard.
- With both FC Flex IO modules present in the MXL or I/O Aggregator switches, the power supply requirement and maximum thermal output are the same as these parameters needed for the M1000 chassis.
- Each port on the FC Flex IO module contains status indicators to denote the link status and transmission activity. For traffic that is being transmitted, the port LED shows a blinking green light. The Link LED displays solid green when a proper link with the peer is established. If there is no connectivity, the LEDs are not lit
- The MXL and IOA switches continue to operate in FCoE Gateway mode even if connectivity to a ToR switch does not exist.
- The I/O Aggregator examines whether the FC Flex IO module is inserted into the switch. When the FC Flex IO module is present during the boot process, the switch runs in FCoE NPIV gateway mode by default.
- When an FC Flex IO module is present in the I/O Aggregator, the software autoconfigures the DCB settings on the ports that support DCB and does not retrieve these settings from the ToR switch.
- Active fabric manager (AFM) is compatible with FC Flex IO modules.
- All SNMP MIBs that are supported for MXL and IOA switches apply equally for FC Flex IO modules. The interface MIB indicates the FC interface when you install the FC flex IO module. The interface MIB statistical counters compute and display the FC interface metrics.
- When the Dell Networking OS sends FC frames (the initial FLOGI or FLOGO messages), or converts FLOGI to FDISC messages or processes any internally generated FC frames, the software computes and verifies the FC cyclic redundancy check (CRC) value before sending the frame to FC ports.
- Fabric worldwide name (WWN) verification is available for eight FC ports. Single-switching WWN capability is provided when the switch operates in NPIV mode.
- With FC Flex IO modules, you can connect the IOA in Simple MUX mode to a single fabric.
- With FC Flex IO modules on an IOA, the FC port speed is set to auto. The following parameters are automatically configured on the ENode facing and FC ports
 - Description: SAN_FABRIC
 - Fabric-id: 1002
 - Fcoe-vlan: 1002
 - Fc-map: 0x0efc00
 - Fcf-priority: 128
 - Fka-adv-period: 8000mSec
 - Keepalive: enable
 - Vlan priority: 3
- On an IOA, the FCoE virtual local area network (VLAN) is automatically configured.
- With FC Flex IO modules on an IOA, the following DCB maps are applied on all of the ENode facing ports.
 - dcb-map: SAN_DCB_MAP
 - priority-group 0 bandwidth 30 pfc off
 - priority-group 1 bandwidth 30 pfc off

- priority-group 2 bandwidth 40 pfc on
- priority-pgid 0 0 0 2 1 0 0 0
- On I/O Aggregators, uplink failure detection (UFD) is disabled if FC Flex IO module is present to allow server ports to communicate with the FC fabric even when the Ethernet upstream ports are not operationally up.
- Ensure that the NPIV functionality is enabled on the upstream switches that operate as FC switches or FCoE forwarders (FCF) before you connect the FC port of the MXL or I/O Aggregator to these upstream switches.
- While storage traffic traverses through FC Flex IO modules and the Ethernet uplink port-channel status changes (with DCB enabled on an adjacent switch), FCoE traffic is disrupted. This problem does not occur if Ethernet traffic is not involved and only FCoE traffic is transmitted. Also, if DCB on the ToR switch is disabled, traffic disruption does not occur.

Port Numbering for FC Flex IO Modules

Even-numbered ports are at the bottom of the I/O panel and for modules odd-numbered ports are at the top of the I/O panel. When installed in a PowerEdge M1000e Enclosure, the MXL 10/40GbE Switch and Aggregator ports are numbered 33 to 56 from the bottom to the top of the switch. The following port numbering convention applies to the FC Flex IO module:

- In expansion slot 0, the ports are numbered 41 to 44.
- In expansion slot 1, the ports are numbered 49 to 52.

Installing the Optics

The following optical ports are supported on the FC Flex IO module using one of the supported breakout cables:

- 4G or 8G Fibre Channel small form-factor pluggable plus (SFP+) optics module and LC connectors over a distance of 150 meters.
- 4G or 8G Fibre Channel SFP+ optics module and LC connectors over a distance of 4 km.



CAUTION:

Electrostatic discharge (ESD) damage can occur if the components are mishandled. Always wear an ESD-preventive wrist or heel ground strap when handling the FC Flex IO module and its components.



WARNING: When working with optical fibres, follow all the warning labels and always wear eye protection. Never look directly into the end of a terminated or unterminated fibre or connector as it may cause eye damage.

1.
 - Position the optic so it is in the correct position. The optic has a key that prevents it from being inserted incorrectly.
 - Insert the optic into the port until it gently snaps into place.



NOTE:

1. When you cable the ports, be sure not to interfere with the airflow from the small vent holes above and below the ports.

Processing of Data Traffic

The Dell Networking OS determines the module type that is plugged into the slot. Based on the module type, the software performs the appropriate tasks. The FC Flex IO module encapsulates and decapsulates the FCoE frames. The module directly switches any non-FCoE or non-FIP traffic, and only FCoE frames are processed and transmitted out of the Ethernet network.

When the external device sends FCoE data frames to the switch that contains the FC Flex IO module, the destination MAC address represents one of the Ethernet MAC addresses assigned to FC ports. Based on the destination address, the FCoE header is removed from the incoming packet and the FC frame is transmitted out of the FC port. The flow control mechanism is performed using per-priority flow control to ensure that frame loss does not occur owing to congestion of frames.

Operation of the FIP Application

The NPIV proxy gateway terminates the FIP sessions and responds to FIP messages. The FIP packets are intercepted by the FC Flex IO module and sent to the Dell Networking OS for further analysis. The FIP application responds to the FIP VLAN discovery request from the host based on the configured FCoE VLANs. For every ENode and VN_Port that is logged in, the FIP application responds to keepalive messages for the virtual channel. If the FC link becomes inactive or a logging off of the switch occurs, the FIP engine sends clear virtual link (CVL) messages to the host. The FIP application also responds to solicited advertisements from the end-device. In addition, the FIP application periodically sends advertisement packets to the end-devices for each FCF that is part of the NPIV proxy gateway.

If FC Flex IO modules are installed, the I/O Aggregator does not perform FIP snooping because the FIP frames are terminated on the switch for NPIV operations. However, on MXL Switches, you can configure the switch to operate in FIP Snooping or NPIV mode.

If the MXL 10/40GbE Switch functions in the NPIV mode and you attempt to set the uplink port to be an FCF or a bridge port, a warning message displays and the settings are not saved. On the Aggregator, if the FC module is present, the uplink ports are not automatically set up as FCF or bridge ports. The FC Flex module cannot function as both an NPIV proxy gateway and a FIP snooping bridge at the same time.

Operation of the NPIV Proxy Gateway

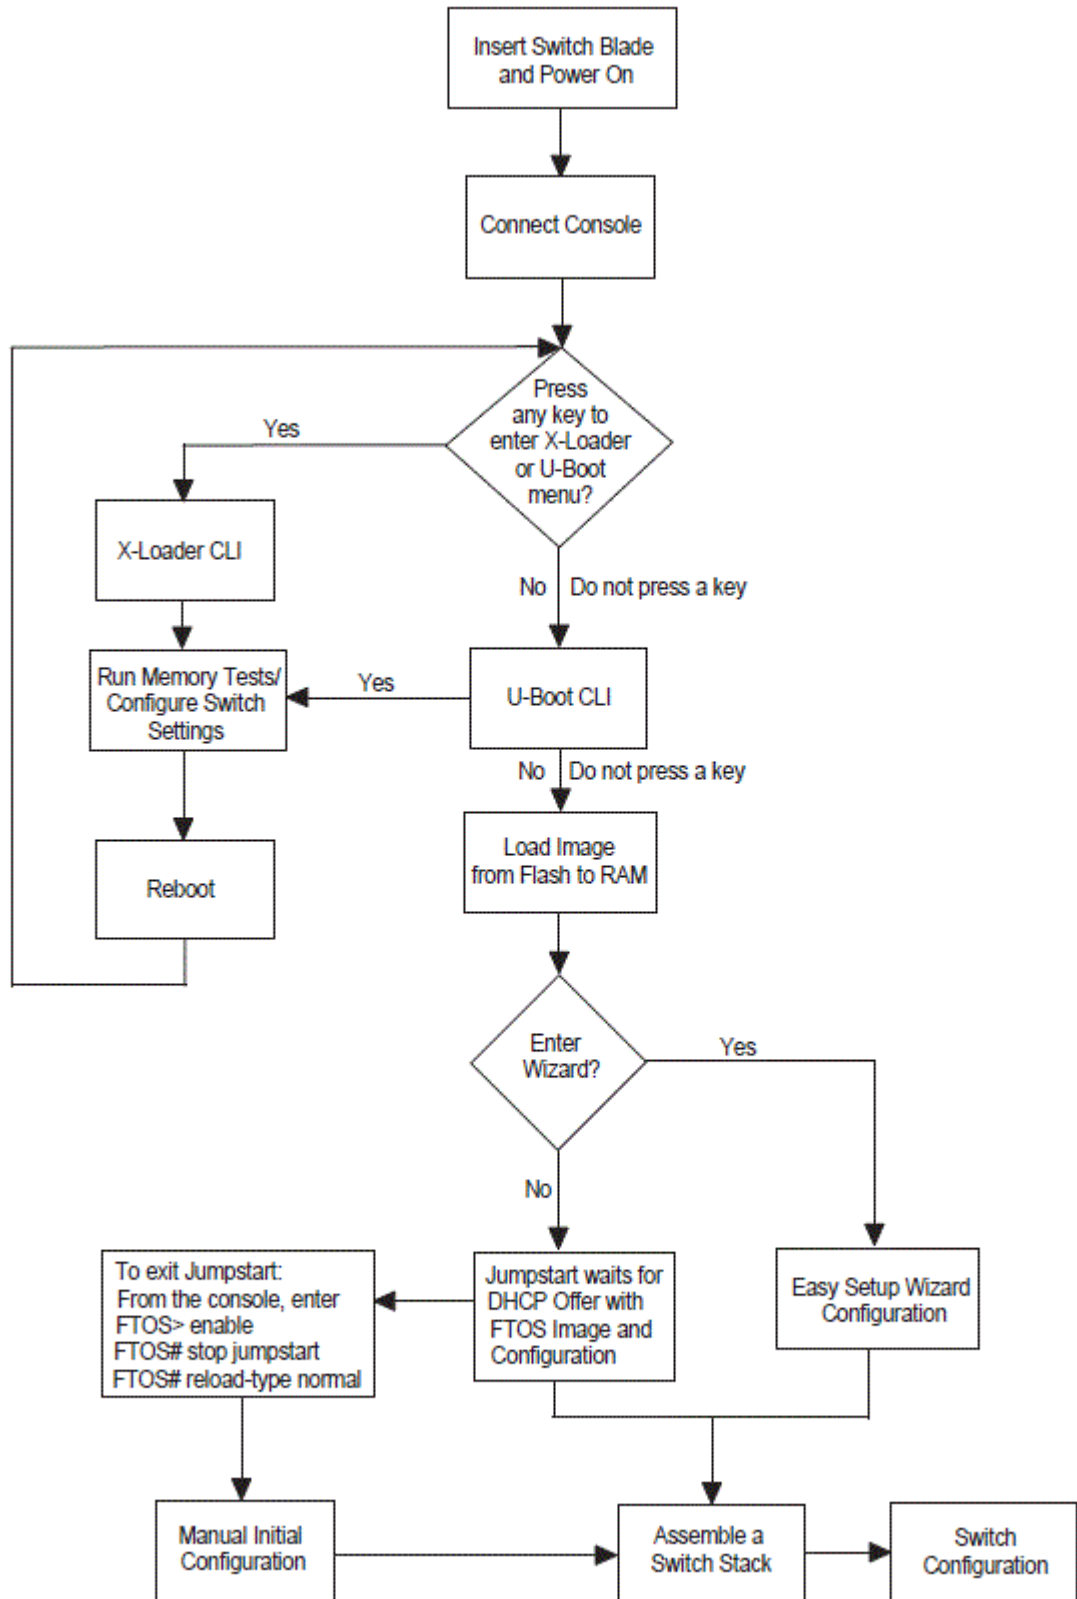
The NPIV application on the FC Flex IO module manages the FC functionalities configured in Dell Networking OS. After the FC link comes up, the gateway sends the initial FLOGI request to the connected switch using the switch and port WWN methods. After a successful login, the NPIV gateway sends a notification to inform the CNA that the FCF is available to log in. The source address of the FIP advertisement and FIP discovery advertisement response contain the MAC address of the FC Flex IO module port. Depending on the number of login sessions on a particular FCF, the NPIV gateway can load-balance the login sessions from ENodes.

The NPIV application performs the FLOGI to FDISC conversion and sends the new FC frame on the associated FC ports. After the external switch responds to the FLOGI request, the NPIV gateway establishes the NPIV session and sends the frame to the FIP application. The FIP application establishes virtual links to convert FCoE FLOGI accept messages into FIP FLOGI accept messages. The corresponding ACL for the accept message is then applied. If a FIP timeout from ENode or VN_PORT occurs, the NPIV application performs the FC fabric logout to the external FC switch. The NPIV application manages the sessions between the FCoE and the FC domain.

Installing and Configuring the Switch

After you unpack the MXL 10/40GbE Switch, refer to the flow chart in the following figure for an overview of the steps you must follow to install the blade and perform the initial configuration.

Installing and Configuring Flowchart for FC Flex IO Modules



To see if a switch is running the latest Dell Networking OS version, use the `show version` command. To download a Dell Networking OS version, go to <http://support.dell.com>.

Installation

Site Preparation

Before installing the switch or switches, make sure that the chosen installation location meets the following site requirements:

- **Clearance** — There is adequate front and rear clearance for operator access. Allow clearance for cabling, power connections, and ventilation.
- **Cabling** — The cabling is routed to avoid sources of electrical noise such as radio transmitters, broadcast amplifiers, power lines, and fluorescent lighting fixtures.
- **Ambient Temperature** — The ambient switch operating temperature range is 10° to 35°C (50° to 95°F).
 1. Decrease the maximum temperature by 1°C (1.8°F) per 300 m (985 ft.) above 900 m (2955 ft.).
 2. **Relative Humidity** — The operating relative humidity is 8 percent to 85 percent (non-condensing) with a maximum humidity gradation of 10 percent per hour.

Unpacking the Switch

Package Contents

When unpacking each switch, make sure that the following items are included:

- One Dell Networking MXL 10/40GbE Switch IO Module
- One USB type A-to-DB-9 female cable
- Getting Started Guide
- Safety and Regulatory Information
- Warranty and Support Information
- Software License Agreement

Unpacking Steps

1. Before unpacking the switch, inspect the container and immediately report any evidence of damage.
2. Place the container on a clean, flat surface and cut all straps securing the container.
3. Open the container or remove the container top.
4. Carefully remove the switch from the container and place it on a secure and clean surface.
5. Remove all packing material.
6. Inspect the product and accessories for damage.

After you insert a Flex IO module into an empty slot, you must reload the I/O Aggregator for the module. If you remove an installed module and insert a different module type, an error message displays to remind you that the slot is configured for a different type of Flex IO module. You must reload the switch to make the Flex IO module operational.

Interconnectivity of FC Flex IO Modules with Cisco MDS Switches

In a network topology that contains Cisco MDS switches, FC Flex IO modules that are plugged into the MXL and I/O Aggregator switches enable interoperation for a robust, effective deployment of the NPIV proxy gateway and FCoE-FC bridging behavior. In an environment that contains FC Flex IO modules and Cisco MDS switches, perform the following steps:

- Insert the FC Flex IO module into any of the optional module slots of the MXL 10/40GBE Switch or the I/O Aggregator Switch and reload the switch.
- When the device is reloaded, NPIV mode is automatically enabled.
- Configure the NPIV-related commands on MXL or I/O Aggregator.

After you perform the preceding procedure, the following operations take place:

- A physical link is established between the FC Flex I/O module and the Cisco MDS switch.
- The FC Flex I/O module sends a proxy FLOGI request to the upstream F_Port of the FC switch or the MDS switch. The F_port accepts the proxy FLOGI request for the FC Flex IO virtual N_Port. The converged network adapters (CNAs) are brought online and the FIP application is run.
- Discovery of the VLAN and FCF MAC addresses is completed.
- The CNA sends a FIP fabric login (FLOGI) request to the FC Flex IO module, which converts FLOGI to FDISC messages or processes any internally generated FC frames and sends these messages to the SAN environment.
- When the FC fabric discovery (FDISC) accept message is received from the SAN side, the FC Flex IO module converts the FDISC message again into an FLOGI accept message and transmits it to the CNA.
- Internal tables of the switch are then programmed to enable the gateway device to forward FCoE traffic directly back and forth between the devices.
- The FC Flex IO module sends an FC or FCoE registered state change notification (RSCN) message to the upstream or downstream devices whenever an error occurs in the appropriate direction.
- An F_Port is a port on an FC switch that connects to an N_Port of an FC device and is called a fabric port.

By default, the NPIV functionality is disabled on the Cisco MDS switch; enable this capability before you connect the FC port of the MXL or I/O Aggregator to these upstream switches.

Data Center Bridging, Fibre Channel over Ethernet, and NPIV Proxy Gateway features are supported on the FC Flex IO modules. For detailed information about these applications and their working, see the corresponding chapters for these applications in this manual.

The following figures illustrate two deployment scenarios of configuring FC Flex IO modules:




Figure 137. Case 1: Deployment Scenario of Configuring FC Flex IO Modules

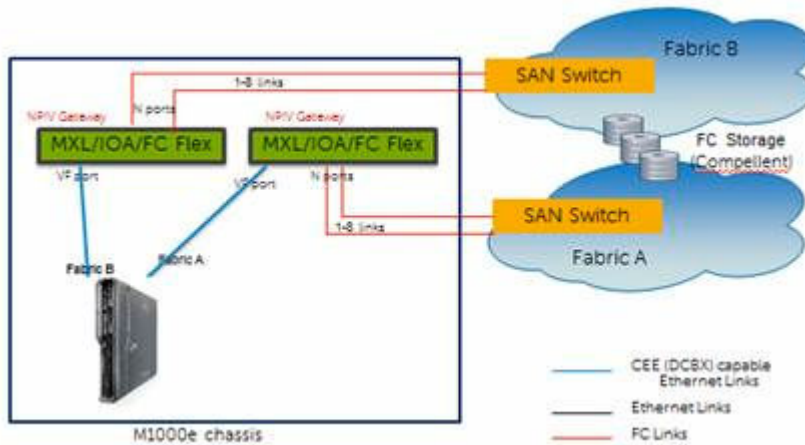


Figure 138. Case 2: Deployment Scenario of Configuring FC Flex IO Modules

Data Center Bridging (DCB)

Data center bridging (DCB) is supported on the FC Flex IO module installed in the MXL 10/40GbE Switch.

Ethernet Enhancements in Data Center Bridging

The following section describes DCB.

- The device supports the following DCB features:
 - Data center bridging exchange protocol (DCBx)
 - Priority-based flow control (PFC)
 - Enhanced transmission selection (ETS)

DCB refers to a set of IEEE Ethernet enhancements that provide data centers with a single, robust, converged network to support multiple traffic types, including local area network (LAN), server, and storage traffic. Through network consolidation, DCB results in reduced operational cost, simplified management, and easy scalability by avoiding the need to deploy separate application-specific networks.

For example, instead of deploying an Ethernet network for LAN traffic, include additional storage area networks (SANs) to ensure lossless Fibre Channel traffic, and a separate InfiniBand network for high-performance inter-processor computing within server clusters, only one DCB-enabled network is required in a data center. The Dell Networking switches that support a unified fabric and consolidate multiple network infrastructures use a single input/output (I/O) device called a converged network adapter (CNA).

A CNA is a computer input/output device that combines the functionality of a host bus adapter (HBA) with a network interface controller (NIC). Multiple adapters on different devices for several traffic types are no longer required.

Data center bridging satisfies the needs of the following types of data center traffic in a unified fabric:

LAN traffic	LAN traffic consists of many flows that are insensitive to latency requirements, while certain applications, such as streaming video, are more sensitive to latency. Ethernet functions as a best-effort network that may drop packets in the case of network congestion. IP networks rely on transport protocols (for example, TCP) for reliable data transmission with the associated cost of greater processing overhead and performance impact.
Storage traffic	Storage traffic based on Fibre Channel media uses the SCSI protocol for data transfer. This traffic typically consists of large data packets with a payload of 2K bytes that cannot recover from frame loss. To successfully transport storage traffic, data center Ethernet must provide no-drop service with lossless links.
InterProcess Communication (IPC) traffic	InterProcess Communication (IPC) traffic within high-performance computing clusters to share information. Server traffic is extremely sensitive to latency requirements.

To ensure lossless delivery and latency-sensitive scheduling of storage and service traffic and I/O convergence of LAN, storage, and server traffic over a unified fabric, IEEE data center bridging adds the following extensions to a classical Ethernet network:

- 802.1Qbb — Priority-based Flow Control (PFC)
- 802.1Qaz — Enhanced Transmission Selection (ETS)
- 802.1Qau — Congestion Notification

- Data Center Bridging Exchange (DCBx) protocol

NOTE: In the Dell Networking OS version 8.3.12.0, only the PFC, ETS, and DCBx features are supported in data center bridging.

Priority-Based Flow Control

In a data center network, priority-based flow control (PFC) manages large bursts of one traffic type in multiprotocol links so that it does not affect other traffic types and no frames are lost due to congestion.

When PFC detects congestion on a queue for a specified priority, it sends a pause frame for the 802.1p priority traffic to the transmitting device. In this way, PFC ensures that PFC-enabled priority traffic is not dropped by the switch.

PFC enhances the existing 802.3x pause and 802.1p priority capabilities to enable flow control based on 802.1p priorities (classes of service). Instead of stopping all traffic on a link (as performed by the traditional Ethernet pause mechanism), PFC pauses traffic on a link according to the 802.1p priority set on a traffic type. You can create lossless flows for storage and server traffic while allowing for loss in case of LAN traffic congestion on the same physical interface.

The following illustration shows how PFC handles traffic congestion by pausing the transmission of incoming traffic with dot1p priority 3.




Figure 139. Priority-Based Flow Control

In the system, PFC is implemented as follows:

- PFC is supported on specified 802.1p priority traffic (dot1p 0 to 7) and is configured per interface. However, only two lossless queues are supported on an interface: one for Fibre Channel over Ethernet (FCoE) converged traffic and one for Internet Small Computer System Interface (iSCSI) storage traffic. Configure the same lossless queues on all ports.
- PFC delay constraints place an upper limit on the transmit time of a queue after receiving a message to pause a specified priority.
- By default, PFC is enabled on an interface with no dot1p priorities configured. You can configure the PFC priorities if the switch negotiates with a remote peer using DCBX.
- During DCBX negotiation with a remote peer:
 - If the negotiation succeeds and the port is in DCBX Willing mode to receive a peer configuration, PFC parameters from the peer are used to configure PFC priorities on the port. If you enable the link-level flow control mechanism on the interface, DCBX negotiation with a peer is not performed.
 - If the negotiation fails and PFC is enabled on the port, any user-configured PFC input policies are applied. If no PFC input policy has been previously applied, the PFC default setting is used (no

priorities configured). If you do not enable PFC on an interface, you can enable the 802.3x link-level pause function. By default, the link-level pause is disabled.

- PFC supports buffering to receive data that continues to arrive on an interface while the remote system reacts to the PFC operation.
- PFC uses the DCB MIB IEEE802.1azd2.5 and the PFC MIB IEEE802.1bb-d2.2.

Enhanced Transmission Selection

Enhanced transmission selection (ETS) supports optimized bandwidth allocation between traffic types in multiprotocol (Ethernet, FCoE, SCSI) links.

ETS allows you to divide traffic according to its 802.1p priority into different priority groups (traffic classes) and configure bandwidth allocation and queue scheduling for each group to ensure that each traffic type is correctly prioritized and receives its required bandwidth. For example, you can prioritize low-latency storage or server cluster traffic in a traffic class to receive more bandwidth and restrict best-effort LAN traffic assigned to a different traffic class.

Although you can configure strict-priority queue scheduling for a priority group, ETS introduces flexibility that allows the bandwidth allocated to each priority group to be dynamically managed according to the amount of LAN, storage, and server traffic in a flow. Unused bandwidth is dynamically allocated to prioritized priority groups. Traffic is queued according to its 802.1p priority assignment, while flexible bandwidth allocation and the configured queue-scheduling for a priority group is supported.

The following figure shows how ETS allows you to allocate bandwidth when different traffic types are classed according to 802.1p priority and mapped to priority groups.




Figure 140. Enhanced Transmission Selection

The following table lists the traffic groupings ETS uses to select multiprotocol traffic for transmission.

Table 78. ETS Traffic Groupings

Traffic Groupings	Description
Priority group	A group of 802.1p priorities used for bandwidth allocation and queue scheduling. All 802.1p priority

Traffic Groupings	Description
Group ID	traffic in a group must have the same traffic handling requirements for latency and frame loss. A 4-bit identifier assigned to each priority group. The range is from 0 to 7.
Group bandwidth	Percentage of available bandwidth allocated to a priority group.
Group transmission selection algorithm (TSA)	Type of queue scheduling a priority group uses.

In the Dell Networking OS, ETS is implemented as follows:

- ETS supports groups of 802.1p priorities that have:
 - PFC enabled or disabled
 - No bandwidth limit or no ETS processing
- Bandwidth allocated by the ETS algorithm is made available after strict-priority groups are serviced. If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups.
- For ETS traffic selection, an algorithm is applied to priority groups using:
 - Strict priority shaping
 - ETS shaping
- ETS uses the DCB MIB IEEE 802.1azd2.5.

Configuring DCB Maps and its Attributes

This topic contains the following sections that describe how to configure a DCB map, apply the configured DCB map to a port, configure PFC without a DCB map, and configure lossless queues.

DCB Map: Configuration Procedure

A DCB map consists of PFC and ETS parameters. By default, PFC is not enabled on any 802.1p priority and ETS allocates equal bandwidth to each priority. To configure user-defined PFC and ETS settings, you must create a DCB map.

Step	Task	Command	Command Mode
1	Enter global configuration mode to create a DCB map or edit PFC and ETS settings.	<code>dcb-map name</code>	CONFIGURATION
2	Configure the PFC setting (on or off) and the ETS bandwidth percentage allocated to traffic in each priority group, or whether the priority group traffic should be handled with strict priority scheduling. You can enable PFC on a maximum of two priority queues on an interface. Enabling PFC for dot1p priorities makes the corresponding port queue lossless. The sum of all allocated bandwidth percentages in all groups in the DCB map must be 100%. Strict-priority traffic is serviced first. Afterwards, bandwidth allocated to other	<code>priority-group group_num {bandwidth percentage strict-priority} pfc {on off}</code>	DCB MAP

Step	Task	Command	Command Mode
	<p>priority groups is made available and allocated according to the specified percentages. If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups.</p> <p>Example: priority-group 0 bandwidth 60 pfc off priority-group 1 bandwidth 20 pfc on priority-group 2 bandwidth 20 pfc on priority-group 4 strict-priority pfc off</p> <p>Repeat this step to configure PFC and ETS traffic handling for each priority group.</p>		
3	<p>Specify the dot1p priority-to-priority group mapping for each priority. Priority-group range: 0 to 7. All priorities that map to the same queue must be in the same priority group.</p> <p>Leave a space between each priority group number. For example: priority-pgid 0 0 0 1 2 4 4 4 in which priority group 0 maps to dot1p priorities 0, 1, and 2; priority group 1 maps to dot1p priority 3; priority group 2 maps to dot1p priority 4; priority group 4 maps to dot1p priorities 5, 6, and 7.</p>	<pre>priority-pgid dot1p0_group_num dot1p1_group_num dot1p2_group_num dot1p3_group_num dot1p4_group_num dot1p5_group_num dot1p6_group_num dot1p7_group_num</pre>	DCB MAP

Important Points to Remember

- If you remove a dot1p priority-to-priority group mapping from a DCB map (`no priority pgid` command), the PFC and ETS parameters revert to their default values on the interfaces on which the DCB map is applied. By default, PFC is not applied on specific 802.1p priorities; ETS assigns equal bandwidth to each 802.1p priority.

As a result, PFC and lossless port queues are disabled on 802.1p priorities, and all priorities are mapped to the same priority queue and equally share the port bandwidth.

- To change the ETS bandwidth allocation configured for a priority group in a DCB map, do not modify the existing DCB map configuration. Instead, first create a new DCB map with the desired PFC and ETS settings, and apply the new map to the interfaces to override the previous DCB map settings. Then, delete the original dot1p priority-priority group mapping.

If you delete the dot1p priority-priority group mapping (`no priority pgid` command) before you apply the new DCB map, the default PFC and ETS parameters are applied on the interfaces. This change may create a DCB mismatch with peer DCB devices and interrupt network operation.

Applying a DCB Map on a Port

To apply a DCB map to an Ethernet port, follow these steps:

Step	Task	Command	Command Mode
1	Enter interface configuration mode on an Ethernet port.	<code>interface</code> <code>{tengigabitEthernet slot/</code>	CONFIGURATION

Step	Task	Command	Command Mode
		<i>port fortygigabitEthernet slot/port}</i>	
2	Apply the DCB map on the Ethernet port to configure it with the PFC and ETS settings in the map; for example: Dell# interface tengigabitEthernet 0/0 Dell(config-if-te-0/0)# dcb-map SAN_A_dcb_map1 Repeat Steps 1 and 2 to apply a DCB map to more than one port. You cannot apply a DCB map on an interface that has been already configured for PFC using the <code>pfcc priority</code> command or which is already configured for lossless queues (<code>pfcc no-drop queues</code> command).	<i>dcb-map name</i>	INTERFACE

Configuring PFC without a DCB Map

In a network topology that uses the default ETS bandwidth allocation (assigns equal bandwidth to each priority), you can also enable PFC for specific dot1p-priorities on individual interfaces without using a DCB map. This type of DCB configuration is useful on interfaces that require PFC for lossless traffic, but do not transmit converged Ethernet traffic.

Step	Task	Command	Command Mode
1	Enter interface configuration mode on an Ethernet port.	<i>interface {tengigabitEthernet slot/port fortygigabitEthernet slot/port}</i>	CONFIGURATION
2	Enable PFC on specified priorities. Range: 0-7. Default: None. Maximum number of lossless queues supported on an Ethernet port: 2. Separate priority values with a comma. Specify a priority range with a dash, for example: <code>pfcc priority 3,5-7</code> 1. You cannot configure PFC using the <code>pfcc priority</code> command on an interface on which a DCB map has been applied or which is already configured for lossless queues (<code>pfcc no-drop queues</code> command).	<i>pfcc priority priority-range</i>	INTERFACE

Configuring Lossless Queues

DCB also supports the manual configuration of lossless queues on an interface after you disable PFC mode in a DCB map and apply the map on the interface. The configuration of no-drop queues provides flexibility for ports on which PFC is not needed, but lossless traffic should egress from the interface.

Lossless traffic egresses out the no-drop queues. Ingress 802.1p traffic from PFC-enabled peers is automatically mapped to the no-drop egress queues.

When configuring lossless queues on a port interface, consider the following points:

- By default, no lossless queues are configured on a port.
- A limit of two lossless queues are supported on a port. If the number of lossless queues configured exceeds the maximum supported limit per port (two), an error message is displayed. You must re-configure the value to a smaller number of queues.
- If you configure lossless queues on an interface that already has a DCB map with PFC enabled (**pfc on**), an error message is displayed.

Step	Task	Command	Command Mode
1	Enter INTERFACE Configuration mode.	<code>interface{tengigabitEthernet slot/port fortygigabitEthernet slot/port}</code>	CONFIGURATION
2	Open a DCB map and enter DCB map configuration mode.	<code>dcb-map name</code>	INTERFACE
3	Disable PFC.	<code>no pfc mode on</code>	DCB MAP
4	Return to interface configuration mode.	<code>exit</code>	DCB MAP
5	Apply the DCB map, created to disable the PFC operation, on the interface	<code>dcb-map {name default}</code>	INTERFACE
6	Configure the port queues that still function as no-drop queues for lossless traffic.	<code>pfc no-drop queue-queue-range</code>	INTERFACE
	The maximum number of lossless queues globally supported on a port is 2.		
	You cannot configure PFC no-drop queues on an interface on which a DCB map with PFC enabled has been applied, or which is already configured for PFC using the <code>pfc priority</code> command.		
	Range: 0-3. Separate queue values with a comma; specify a priority range with a dash; for example: <code>pfc no-drop queues 1,3</code> or <code>pfc no-drop queues 2-3</code> Default: No lossless queues are configured.		

Data Center Bridging Exchange Protocol (DCBx)

DCBx allows a switch to automatically discover DCB-enabled peers and exchange configuration information. PFC and ETS use DCBx to exchange and negotiate parameters with peer devices. DCBx capabilities include:

- Discovery of DCB capabilities on peer-device connections.
- Determination of possible mismatch in DCB configuration on a peer link.
- Configuration of a peer device over a DCB link.

DCBx requires the link layer discovery protocol (LLDP) to provide the path to exchange DCB parameters with peer devices. Exchanged parameters are sent in organizationally specific TLVs in LLDP data units. For more information, refer to [Link Layer Discovery Protocol \(LLDP\)](#). The following LLDP TLVs are supported for DCB parameter exchange:

PFC parameters PFC Configuration TLV and Application Priority Configuration TLV.

ETS parameters ETS Configuration TLV and ETS Recommendation TLV.

Data Center Bridging in a Traffic Flow

The following figure shows how DCB handles a traffic flow on an interface.

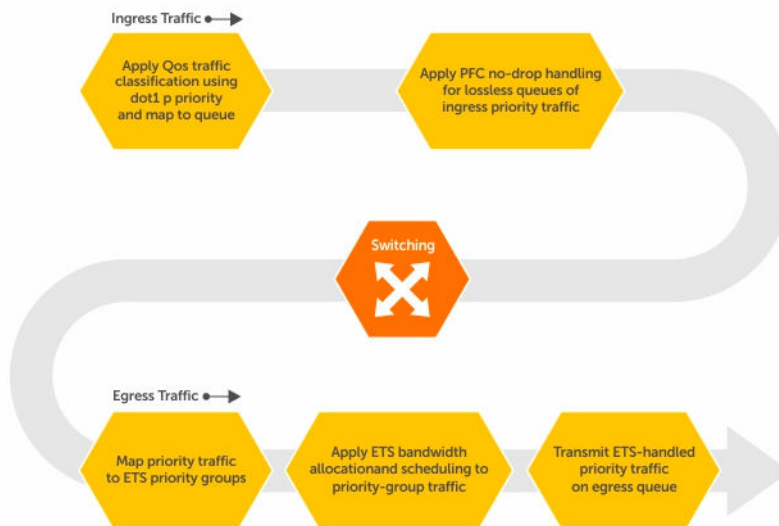


Figure 141. DCB PFC and ETS Traffic Handling

Enabling Data Center Bridging

Data center bridging is enabled by default on an MXL 10/40GbE Switch to support converged enhanced Ethernet (CEE) in a data center network. A prerequisite for configuring DCB:

- Priority-based flow control
- Enhanced transmission selection
- Data center bridging exchange protocol
- FCoE initialization protocol (FIP) snooping

DCB processes virtual local area network (VLAN)-tagged packets and dot1p priority values. Untagged packets are treated with a dot1p priority of 0.

For DCB to operate effectively, you can classify ingress traffic according to its dot1p priority so that it maps to different data queues. The dot1p-queue assignments used are shown in the following table.


On the MXL Switch, by default, DCB is enabled and MMU buffers are reserved to achieve no-drop traffic handling for PFC. Disabling DCB does not release the buffers reserved by default. To utilize reserved buffers for non-DCB applications, you have to explicitly release the buffers (Refer to [Configuring the PFC Buffer in a Switch Stack](#)).

To disable or re-enable DCB on a switch, enter the following commands.

1. Disable DCB.
CONFIGURATION mode

no dcb enable
2. Re-enable DCB.
CONFIGURATION mode

dcb enable

 **NOTE: Dell Networking OS Behavior:** DCB is not supported if you enable link-level flow control on one or more interfaces.


After you disable DCB, if link-level flow control is not automatically enabled on an interface, to enable flow control, manually shut down the interface (the `shutdown` command) and re-enable it (the `no shutdown` command).

QoS dot1p Traffic Classification and Queue Assignment

The following section describes QoS dot1P traffic classification and assignments.


DCB supports PFC, ETS, and DCBx to handle converged Ethernet traffic that is assigned to an egress queue according to the following QoS methods:

- | | |
|---------------------------|---|
| Honor dot1p | You can honor dot1p priorities in ingress traffic at the port or global switch level (refer to Default dot1p to Queue Mapping) using the <code>service-class dynamic dot1p</code> command in INTERFACE configuration mode (refer to Honoring dot1p Values on Ingress Packets). |
| Layer 2 class maps | You can use dot1p priorities to classify traffic in a class map and apply a service policy to an ingress port to map traffic to egress queues (refer to Policy-Based QoS Configurations). |

 **NOTE:** Dell Networking does not recommend mapping all ingress traffic to a single queue when using PFC and ETS. However, Dell Networking does recommend using Ingress traffic classification using the `service-class dynamic dot1p` command (honor dot1p) on all DCB-enabled interfaces. If you use L2 class maps to map dot1p priority traffic to egress queues, take into account the default dot1p-queue assignments in the following table and the maximum number of two lossless queues supported on a port (refer to [Configuring Lossless Queues](#)).

Although the system allows you to change the default dot1p priority-queue assignments (refer to [Setting dot1p Priorities for Incoming Traffic](#)), DCB policies applied to an interface may become invalid if you reconfigure dot1p-queue mapping. If the configured DCB policy remains valid, the change in the dot1p-queue assignment is allowed. For DCB ETS enabled interfaces, traffic destined to queue that is not mapped to any dot1p priority are dropped.

dot1p Value in the Incoming Frame	Egress Queue Assignment
0	0
1	0
2	0
3	1
4	2
5	3
6	3
7	3

 **NOTE:** If you reconfigure the global dot1p-queue mapping, an automatic re-election of the DCBX configuration source port is performed (refer to [Configuration Source Election](#)).

Configure Enhanced Transmission Selection

ETS provides a way to optimize bandwidth allocation to outbound 802.1p classes of converged Ethernet traffic.

Different traffic types have different service needs. Using ETS, you can create groups within an 802.1p priority class to configure different treatment for traffic with different bandwidth, latency, and best-effort needs.

For example, storage traffic is sensitive to frame loss; interprocess communication (IPC) traffic is latency-sensitive. ETS allows different traffic types to coexist without interruption in the same converged link by:

- Allocating a guaranteed share of bandwidth to each priority group.
- Allowing each group to exceed its minimum guaranteed bandwidth if another group is not fully using its allotted bandwidth.

To configure ETS and apply an ETS output policy to an interface, you must:

1. Create a Quality of Service (QoS) output policy with ETS scheduling and bandwidth allocation settings.
2. Create a priority group of 802.1p traffic classes.

3. Configure a DCB output policy in which you associate a priority group with a QoS ETS output policy.
4. Apply the DCB output policy to an interface.

ETS Operation with DCBx

The following section describes DCBx negotiation with peer ETS devices.

In DCBx negotiation with peer ETS devices, ETS configuration is handled as follows:

- ETS TLVs are supported in DCBx versions CIN, CEE, and IEEE2.5.
- The DCBx port-role configurations determine the ETS operational parameters (refer to [Configure a DCBx Operation](#)).
- ETS configurations received from TLVs from a peer are validated.
- If there is a hardware limitation or TLV error:
 - DCBx operation on an ETS port goes down.
 - New ETS configurations are ignored and existing ETS configurations are reset to the previously configured ETS output policy on the port or to the default ETS settings if no ETS output policy was previously applied.
- ETS operates with legacy DCBx versions as follows:
 - In the CEE version, the priority group/traffic class group (TCG) ID 15 represents a non-ETS priority group. Any priority group configured with a scheduler type is treated as a strict-priority group and is given the priority-group (TCG) ID 15.
 - The CIN version supports two types of strict-priority scheduling:
 - * Group strict priority: Use this to increase its bandwidth usage to the bandwidth total of the priority group and allow a single priority flow in a priority group. A single flow in a group can use all the bandwidth allocated to the group.
 - * Link strict priority: Use this to increase to the maximum link bandwidth and allow a flow in any priority group.

CIN supports only the dot1p priority-queue assignment in a priority group. To configure a dot1p priority flow in a priority group to operate with link strict priority, you configure: The dot1p priority for strict-priority scheduling (`strict-priority` command; [Enabling Strict-Priority Queueing](#)).

If you configure only the priority group in an ETS output policy or only the dot1p priority for strict-priority scheduling, the flow is handled with group strict priority.

Configuring Bandwidth Allocation for DCBx CIN

After you apply an ETS output policy to an interface, if the DCBx version used in your data center network is CIN, you may need to configure a QoS output policy to overwrite the default CIN bandwidth allocation. This default setting divides the bandwidth allocated to each port queue equally between the dot1p priority traffic assigned to the queue.

For more information, refer to [Allocating Bandwidth to Queue](#).

To create a QoS output policy that allocates different amounts of bandwidth to the different traffic types/dot1p priorities assigned to a queue and apply the output policy to the interface, follow these steps.

1. Create a QoS output policy.

CONFIGURATION mode

```
Dell (conf) #qos-policy-output test12
```


The maximum 32 alphanumeric characters.

2. Configure the percentage of bandwidth to allocate to the dot1p priority/queue traffic in the associated L2 class map.

QoS OUTPUT POLICY mode

```
Dell(conf-qos-policy-out)#bandwidth-percentage 100
```

The default is **none**.

3. Repeat Step 2 to configure bandwidth percentages for other priority queues on the port.

QoS OUTPUT POLICY mode

```
Dell(conf-qos-policy-out)#bandwidth-percentage 100
```

4. Create a priority group for strict-priority scheduling.

QoS OUTPUT POLICY mode

```
Dell(conf-qos-policy-out)#scheduler strict
```



NOTE: You can not use `scheduler strict` when bandwidth percentage is configured. It displays an error message.

```
Dell(conf-qos-policy-out)#bandwidth-percentage 100
Dell(conf-qos-policy-out)#scheduler strict
% Error: Strict priority scheduler mode is not allowed when bandwidth-
percentage is configured on qos-policy-output profile.
Dell(conf-qos-policy-out)#scheduler strict ?
```

5. Exit QoS Output Policy Configuration mode.

QoS OUTPUT POLICY mode

```
Dell(conf-if-te-0/1)#exit
```

6. Enter INTERFACE Configuration mode.

CONFIGURATION mode

```
Dell(conf-qos-policy-out)#int te 0/1
```

7. Apply the QoS output policy with the bandwidth percentage for specified priority queues to an egress interface.

INTERFACE mode

```
Dell(conf-if-te-0/1)#service-policy output test12
```

Configure a DCBx Operation

DCB devices use data center bridging exchange protocol (DCBx) to exchange configuration information with directly connected peers using the link layer discovery protocol (LLDP) protocol.

DCBx can detect the misconfiguration of a peer DCB device, and optionally, configure peer DCB devices with DCB feature settings to ensure consistent operation in a data center network.

DCBx is a prerequisite for using DCB features, such as priority-based flow control (PFC) and enhanced traffic selection (ETS), to exchange link-level configurations in a converged Ethernet environment. DCBx is also deployed in topologies that support lossless operation for FCoE or iSCSI traffic. In these scenarios, all network devices are DCBx-enabled (DCBx is enabled end-to-end). For more information about how these features are implemented and used, refer to:

- [Configuring Priority-Based Flow Control](#)
- [Configure Enhanced Transmission Selection](#)
- [Configuring FIP Snooping](#)

DCBx supports the following versions: CIN, CEE, and IEEE2.5.

Prerequisite: For DCBx, enable LLDP on all DCB devices.

DCBx Operation

DCBx performs the following operations:

- Discovers DCB configuration (such as PFC and ETS) in a peer device.
- Detects DCB mis-configuration in a peer device; that is, when DCB features are not compatibly configured on a peer device and the local switch. Mis-configuration detection is feature-specific because some DCB features support asymmetric configuration.
- Reconfigures a peer device with the DCB configuration from its configuration source if the peer device is willing to accept configuration.
- Accepts the DCB configuration from a peer if a DCBx port is in “willing” mode to accept a peer’s DCB settings and then internally propagates the received DCB configuration to its peer ports.

DCBx Port Roles

To enable the auto-configuration of DCBx-enabled ports and propagate DCB configurations learned from peer DCBx devices internally to other switch ports, use the following DCBx port roles.

Auto-upstream

The port advertises its own configuration to DCBx peers and receives its configuration from DCBX peers (ToR or FCF device). The port also propagates its configuration to other ports on the switch.

The first auto-upstream that is capable of receiving a peer configuration is elected as the *configuration source*. The elected configuration source then internally propagates the configuration to other auto-upstream and auto-downstream ports. A port that receives an internally propagated configuration overwrites its local configuration with the new parameter values.

When an auto-upstream port (besides the configuration source) receives and overwrites its configuration with internally propagated information, one of the following actions is taken:

- If the peer configuration received is compatible with the internally propagated port configuration, the link with the DCBx peer is enabled.
- If the received peer configuration is not compatible with the currently configured port configuration, the link with the DCBX peer port is disabled and a syslog message for an incompatible configuration is generated. The network administrator must then reconfigure the peer device so that it advertises a compatible DCB configuration.

The configuration received from a DCBX peer or from an internally propagated configuration is not stored in the switch’s running configuration.

On a DCBX port in an auto-upstream role, the PFC and application priority TLVs are enabled. ETS recommend TLVs are disabled and ETS configuration TLVs are enabled.

Auto-downstream

The port advertises its own configuration to DCBx peers but is *not willing* to receive remote peer configuration. The port always accepts internally propagated configurations from a configuration source. An auto-downstream port that receives an internally propagated configuration overwrites its local configuration with the new parameter values.

When an auto-downstream port receives and overwrites its configuration with internally propagated information, one of the following actions is taken:

- If the peer configuration received is compatible with the internally propagated port configuration, the link with the DCBx peer is enabled.
- If the received peer configuration is not compatible with the currently configured port configuration, the link with the DCBX peer port is disabled and a syslog message for an incompatible configuration is generated. The network administrator must then reconfigure the peer device so that it advertises a compatible DCB configuration.

The internally propagated configuration is not stored in the switch's running configuration. On a DCBX port in an auto-downstream role, all PFC, application priority, ETS recommend, and ETS configuration TLVs are enabled.

Configuration source

The port is configured to serve as a source of configuration information on the switch. Peer DCB configurations received on the port are propagated to other DCBx auto-configured ports. If the peer configuration is compatible with a port configuration, DCBx is enabled on the port.

On a configuration-source port, the link with a DCBx peer is enabled when the port receives a DCB configuration that can be internally propagated to other auto-configured ports.

The configuration received from a DCBX peer is not stored in the switch's running configuration.

On a DCBX port that is the configuration source, all PFC and application priority TLVs are enabled. ETS recommend TLVs are disabled and ETS configuration TLVs are enabled.

Manual

The port is configured to operate only with administrator-configured settings and does not auto-configure with DCB settings received from a DCBx peer or from an internally propagated configuration from the configuration source. If you enable DCBx, ports in Manual mode advertise their configurations to peer devices but do not accept or propagate internal or external configurations. Unlike other user-configured ports, the configuration of DCBx ports in Manual mode is saved in the running configuration.

On a DCBx port in a manual role, all PFC, application priority, ETS recommend, and ETS configuration TLVs are enabled.

The default for the DCBx port role is **manual**.



NOTE: On a DCBx port, application priority TLV advertisements are handled as follows:

- The application priority TLV is transmitted only if the priorities in the advertisement match the configured PFC priorities on the port.
- On auto-upstream and auto-downstream ports:
 - If a configuration source is elected, the ports send an application priority TLV based on the application priority TLV received on the configuration-source port. When an application priority TLV is received on the configuration-source port, the auto-upstream and auto-downstream ports use the internally propagated PFC priorities to match against the received application priority. Otherwise, these ports use their locally configured PFC priorities in application priority TLVs.
 - If no configuration source is configured, auto-upstream and auto-downstream ports check to see that the locally configured PFC priorities match the priorities in a received application priority TLV.
- On manual ports, an application priority TLV is advertised only if the priorities in the TLV match the PFC priorities configured on the port.

DCB Configuration Exchange

The DCBx protocol supports the exchange and propagation of configuration information for the enhanced transmission selection (ETS) and priority-based flow control (PFC) DCB features.

DCBx uses the following methods to exchange DCB configuration parameters:

- | | |
|-------------------|---|
| Asymmetric | DCB parameters are exchanged between a DCBx-enabled port and a peer port without requiring that a peer port and the local port use the same configured values for the configurations to be compatible. For example, ETS uses an asymmetric exchange of parameters between DCBx peers. |
| Symmetric | DCB parameters are exchanged between a DCBx-enabled port and a peer port but requires that each configured parameter value be the same for the configurations in order to be compatible. For example, PFC uses an symmetric exchange of parameters between DCBx peers. |

Configuration Source Election

When an auto-upstream or auto-downstream port receives a DCB configuration from a peer, the port first checks to see if there is an active configuration source on the switch.


- If a configuration source already exists, the received peer configuration is checked against the local port configuration. If the received configuration is compatible, the DCBx marks the port as DCBx-enabled. If the configuration received from the peer is not compatible, a warning message is logged and the DCBx frame error counter is incremented. Although DCBx is operationally disabled, the port keeps the peer link up and continues to exchange DCBx packets. If a compatible peer configuration is later received, DCBx is enabled on the port.
- If there is no configuration source, a port may elect itself as the configuration source. A port may become the configuration source if the following conditions exist:
 - No other port is the configuration source.
 - The port role is auto-upstream.
 - The port is enabled with link up and DCBx enabled.
 - The port has performed a DCBx exchange with a DCBx peer.
 - The switch is capable of supporting the received DCB configuration values through either a symmetric or asymmetric parameter exchange.

A newly elected configuration source propagates configuration changes received from a peer to the other auto-configuration ports. Ports receiving auto-configuration information from the configuration source ignore their current settings and use the configuration source information.

Propagation of DCB Information

When an auto-upstream or auto-downstream port receives a DCB configuration from a peer, the port acts as a DCBx client and checks if a DCBx configuration source exists on the switch.

- If a configuration source is found, the received configuration is checked against the currently configured values that are internally propagated by the configuration source. If the local configuration is compatible with the received configuration, the port is enabled for DCBx operation and synchronization.
- If the configuration received from the peer is not compatible with the internally propagated configuration used by the configuration source, the port is disabled as a client for DCBx operation and synchronization and a syslog error message is generated. The port keeps the peer link up and continues to exchange DCBx packets. If a compatible configuration is later received from the peer, the port is enabled for DCBx operation.

 **NOTE:** DCB configurations internally propagated from a configuration source do not overwrite the configuration on a DCBx port in a manual role. When a configuration source is elected, all auto-upstream ports other than the configuration source are marked as *willing disabled*. The internally propagated DCB configuration is refreshed on all auto-configuration ports and each port may begin configuration negotiation with a DCBx peer again.


Auto-Detection and Manual Configuration of the DCBx Version

When operating in Auto-Detection mode (the `DCBx version auto` command), a DCBx port automatically detects the DCBx version on a peer port. Legacy CIN and CEE versions are supported in addition to the standard IEEE version 2.5 DCBx.

A DCBx port detects a peer version after receiving a valid frame for that version. The local DCBx port reconfigures to operate with the peer version and maintains the peer version on the link until one of the following conditions occurs:

- The switch reboots.
- The link is reset (goes down and up).
- User-configured CLI commands require the version negotiation to restart.
- The peer times out.
- Multiple peers are detected on the link.

If you configure a DCBx port to operate with a specific version (the `DCBx version {cee | cin | ieee-v2.5}` command in the [Configuring DCBx](#)), DCBx operations are performed according to the configured version, including fast and slow transmit timers and message formats. If a DCBx frame with a different version is received, a syslog message is generated and the peer version is recorded in the peer status table. If the frame cannot be processed, it is discarded and the discard counter is incremented.

 **NOTE:** Because DCBx TLV processing is best effort, it is possible that CIN frames may be processed when DCBx is configured to operate in CEE mode and vice versa. In this case, the unrecognized TLVs cause the unrecognized TLV counter to increment, but the frame is processed and is not discarded.

Legacy DCBx (CIN and CEE) supports the DCBx control state machine that is defined to maintain the sequence number and acknowledge the number sent in the DCBx control TLVs.

DCBx Example

The following figure shows how DCBx is used on an MXL Switch installed in a PowerEdge M1000e chassis in which servers are also installed.

The external 40GbE ports on the base module (ports 33 and 37) of two switches are used for uplinks configured as DCBx auto-upstream ports. The MXL switch is connected to third-party, top-of-rack (ToR) switches through 40GbE uplinks. The ToR switches are part of a Fibre Channel storage network. The internal ports (ports 1-32) connected to the 10GbE backplane are configured as auto-downstream ports.

On the MXL switch, PFC and ETS use DCBx to exchange link-level configuration with DCBx peer devices.

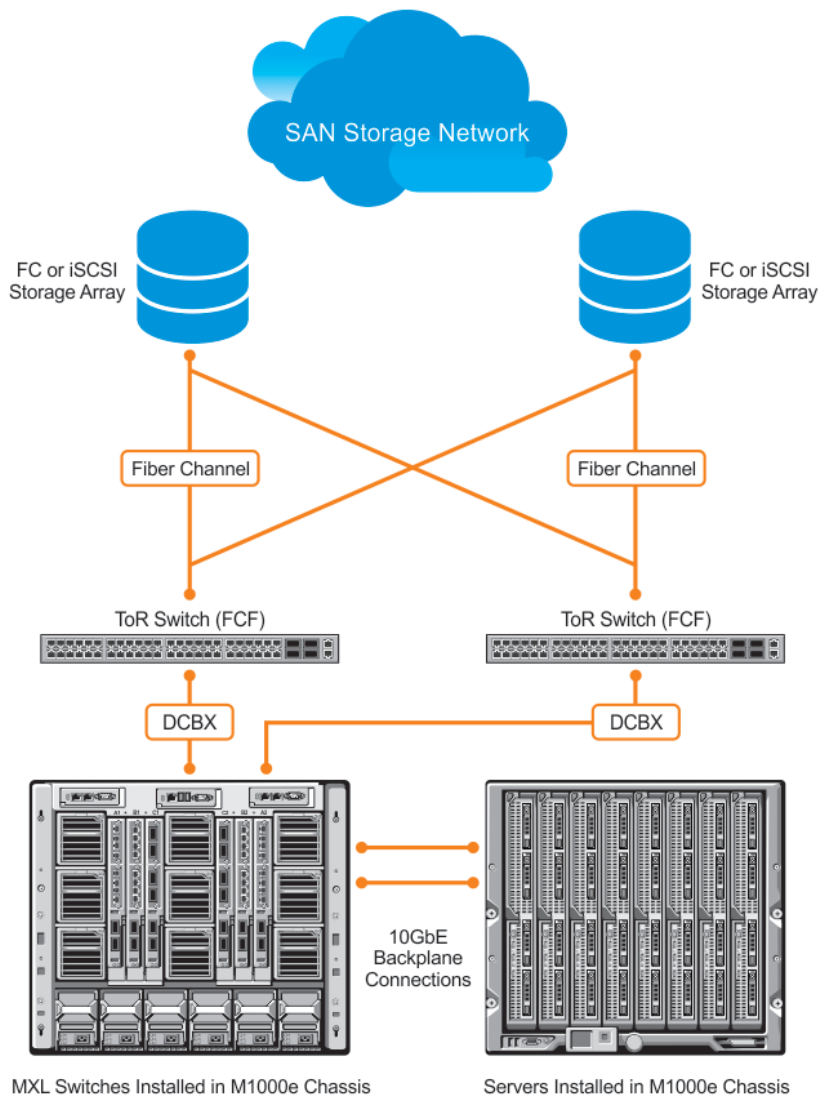


Figure 142. DCBx Sample Topology

DCBx Prerequisites and Restrictions

The following prerequisites and restrictions apply when you configure DCBx operation on a port:

- For DCBx, on a port interface, enable LLDP in both Send (TX) and Receive (RX) mode (the `protocol lldp mode` command; refer to the example in [CONFIGURATION versus INTERFACE Configurations](#) in the [Link Layer Discovery Protocol \(LLDP\)](#) chapter). If multiple DCBx peer ports are detected on a local DCBx interface, LLDP is shut down.
- The CIN version of DCBx supports only PFC, ETS, and FCOE; it does not support iSCSI, backward congestion management (BCN), logical link down (LLDF), and network interface virtualization (NIV).

Configuring DCBx

To configure DCBx, follow these steps.

For DCBx, to advertise DCBx TLVs to peers, enable LLDP. For more information, refer to [Link Layer Discovery Protocol \(LLDP\)](#).

Configure DCBx operation at the interface level on a switch or globally on the switch. To configure an MXL switch for DCBx operation in a data center network, you must:

1. Configure ToR- and FCF-facing interfaces as auto-upstream ports.
2. Configure server-facing interfaces as auto-downstream ports.
3. Configure a port to operate in a configuration-source role.
4. Configure ports to operate in a manual role.

1. Enter INTERFACE Configuration mode.

```
CONFIGURATION mode
```

```
interface type slot/port
```

2. Enter LLDP Configuration mode to enable DCBx operation.

```
INTERFACE mode
```

```
[no] protocol lldp
```

3. Configure the DCBx version used on the interface, where: `auto` configures the port to operate using the DCBx version received from a peer.

```
PROTOCOL LLDP mode
```

```
[no] DCBx version {auto | cee | cin | ieee-v2.5}
```

- `cee`: configures the port to use CEE (Intel 1.01).
- `cin`: configures the port to use Cisco-Intel-Nuova (DCBx 1.0).
- `ieee-v2.5`: configures the port to use IEEE 802.1Qaz (Draft 2.5).

The default is **Auto**.

4. Configure the DCBx port role the interface uses to exchange DCB information.

```
PROTOCOL LLDP mode
```

```
[no] DCBx port-role {config-source | auto-downstream | auto-upstream | manual}
```

- `auto-upstream`: configures the port to receive a peer configuration. The configuration source is elected from auto-upstream ports.
- `auto-downstream`: configures the port to accept the internally propagated DCB configuration from a configuration source.
- `config-source`: configures the port to serve as the configuration source on the switch.
- `manual`: configures the port to operate only on administer-configured DCB parameters. The port does not accept a DCB configuration received from a peer or a local configuration source.


The default is **Manual**.

5. **On manual ports only:** Configure the PFC and ETS TLVs advertised to DCBx peers.
PROTOCOL LLDP mode

```
[no] advertise DCBx-tlv {ets-conf | ets-reco | pfc} [ets-conf | ets-reco | pfc] [ets-conf | ets-reco | pfc]
```

- `ets-conf`: enables the advertisement of ETS Configuration TLVs.
- `ets-reco`: enables the advertisement of ETS Recommend TLVs.
- `pfc enables`: the advertisement of PFC TLVs.

The default is All PFC and ETS TLVs are advertised.

 **NOTE:** You can configure the transmission of more than one TLV type at a time; for example, `advertise DCBx-tlv ets-conf ets-reco`. You can enable ETS recommend TLVs (`ets-reco`) only if you enable ETS configuration TLVs (`ets-conf`).


To disable TLV transmission, use the `no` form of the command; for example, `no advertise DCBx-tlv pfc ets-reco`.

6. **On manual ports only:** Configure the Application Priority TLVs advertised on the interface to DCBx peers.
PROTOCOL LLDP mode

```
[no] advertise DCBx-appln-tlv {fcoe | iscsi}
```

- `fcoe`: enables the advertisement of FCoE in Application Priority TLVs.
- `iscsi`: enables the advertisement of iSCSI in Application Priority TLVs.

The default is Application Priority TLVs are enabled to advertise FCoE and iSCSI.

 **NOTE:** To disable TLV transmission, use the `no` form of the command; for example, `no advertise DCBx-appln-tlv iscsi`.

For information about how to use FCoE and iSCSI, refer to [Fibre Channel over Ethernet](#) and [iSCSI Optimization](#).

To verify the DCBx configuration on a port, use the `show interface DCBx detail` command.

Configuring DCBx Globally on the Switch

To globally configure the DCBx operation on a switch, follow these steps.

1. Enter Global Configuration mode.
EXEC PRIVILEGE mode

configure

2. Enter LLDP Configuration mode to enable DCBx operation.

CONFIGURATION mode

```
[no] protocol lldp
```

3. Configure the DCBx version used on all interfaces not already configured to exchange DCB information.

PROTOCOL LLDP mode

```
[no] DCBx version {auto | cee | cin | ieee-v2.5}
```

- `auto`: configures all ports to operate using the DCBx version received from a peer.
- `cee`: configures a port to use CEE (Intel 1.01). `cin` configures a port to use Cisco-Intel-Nuova (DCBx 1.0).
- `ieee-v2.5`: configures a port to use IEEE 802.1Qaz (Draft 2.5).

The default is **Auto**.



NOTE: To configure the DCBx port role the interfaces use to exchange DCB information, use the `DCBx port-role` command in INTERFACE Configuration mode (Step 3).

4. Configure the PFC and ETS TLVs that advertise on unconfigured interfaces with a manual port-role.

PROTOCOL LLDP mode

```
[no] advertise DCBx-tlv {ets-conf | ets-reco | pfc} [ets-conf | ets-reco | pfc] [ets-conf | ets-reco | pfc]
```

- `ets-conf`: enables transmission of ETS Configuration TLVs.
- `ets-reco`: enables transmission of ETS Recommend TLVs.
- `pfc`: enables transmission of PFC TLVs.



NOTE: You can configure the transmission of more than one TLV type at a time. You can only enable ETS recommend TLVs (`ets-reco`) if you enable ETS configuration TLVs (`ets-conf`). To disable TLV transmission, use the `no` form of the command; for example, `no advertise DCBx-tlv pfc ets-reco`.

The default is All TLV types are enabled.

5. Configure the Application Priority TLVs that advertise on unconfigured interfaces with a manual port-role.

PROTOCOL LLDP mode

```
[no] advertise DCBx-appln-tlv {fcoe | iscsi}
```

- `fcoe`: enables the advertisement of FCoE in Application Priority TLVs.
- `iscsi`: enables the advertisement of iSCSI in Application Priority TLVs.

The default is Application Priority TLVs are enabled and advertise FCoE and iSCSI.



NOTE: To disable TLV transmission, use the `no` form of the command; for example, `no advertise DCBx-appln-tlv iscsi`.

For information about how to use FCoE and iSCSI, refer to [Fibre Channel over Ethernet](#) and [iSCSI Optimization](#).

6. Configure the FCoE priority advertised for the FCoE protocol in Application Priority TLVs.
PROTOCOL LLDP mode

```
[no] fcoe priority-bits priority-bitmap
```

The priority-bitmap range is from 1 to FF.

The default is **0x8**.

7. Configure the iSCSI priority advertised for the iSCSI protocol in Application Priority TLVs.
PROTOCOL LLDP mode

```
[no] iscsi priority-bits priority-bitmap
```

The priority-bitmap range is from 1 to FF.

The default is **0x10**.

DCBx Error Messages

The following syslog messages appear when an error in DCBx operation occurs.

LLDP_MULTIPLE_PEER_DETECTED: DCBx is operationally disabled after detecting more than one DCBx peer on the port interface.

LLDP_PEER_AGE_OUT: DCBx is disabled as a result of LLDP timing out on a DCBx peer interface.

DSM_DCBx_PEER_VERSION_CONFLICT: A local port expected to receive the IEEE, CIN, or CEE version in a DCBx TLV from a remote peer but received a different, conflicting DCBx version.

DSM_DCBx_PFC_PARAMETERS_MATCH and DSM_DCBx_PFC_PARAMETERS_MISMATCH: A local DCBx port received a compatible (match) or incompatible (mismatch) PFC configuration from a peer.

DSM_DCBx_ETS_PARAMETERS_MATCH and DSM_DCBx_ETS_PARAMETERS_MISMATCH: A local DCBx port received a compatible (match) or incompatible (mismatch) ETS configuration from a peer.

LLDP_UNRECOGNISED_DCBx_TLV_RECEIVED: A local DCBx port received an unrecognized DCBx TLV from a peer.

Debugging DCBx on an Interface

To enable DCBx debug traces for all or a specific control paths, use the following command.

- Enable DCBx debugging.
EXEC PRIVILEGE mode

```
debug DCBx {all | auto-detect-timer | config-exchng | fail | mgmt | resource  
| sem | tlv}
```

- all: enables all DCBx debugging operations.

- `auto-detect-timer`: enables traces for DCBx auto-detect timers.
- `config-exchng`: enables traces for DCBx configuration exchanges.
- `fail`: enables traces for DCBx failures.
- `mgmt`: enables traces for DCBx management frames.
- `resource`: enables traces for DCBx system resource frames.
- `sem`: enables traces for the DCBx state machine.
- `tlv`: enables traces for DCBx TLVs.

Verifying the DCB Configuration

To display DCB configurations, use the following `show` commands.

Table 79. Displaying DCB Configurations

Command	Output
<code>show dot1p-queue mapping</code>	Displays the current 802.1p priority-queue mapping.
<code>show dcb [stack-unit <i>unit-number</i>]</code>	Displays the data center bridging status, number of PFC-enabled ports, and number of PFC-enabled queues. On the master switch in a stack, you can specify a stack-unit number. The range is from 0 to 5.
<code>show qos priority-groups</code>	Displays the ETS priority groups configured on the switch, including the 802.1p priority classes and ID of each group.
<code>show interface <i>port-type slot/port</i> pfc {summary detail}</code>	Displays the PFC configuration applied to ingress traffic on an interface, including priorities and link delay. To clear PFC TLV counters, use the <code>clear pfc counters interface <i>port-type slot/port</i></code> command.
<code>show interface <i>port-type slot/port</i> pfc statistics</code>	Displays counters for the PFC frames received and transmitted (by dot1p priority class) on an interface.
<code>show interface <i>port-type slot/port</i> ets {summary detail}</code>	Displays the ETS configuration applied to egress traffic on an interface, including priority groups with priorities and bandwidth allocation. To clear ETS TLV counters, enter the <code>clear ets counters interface <i>port-type slot/port</i></code> command.

Example of the show dot1p-queue mapping Command

Example of the show dcb Command

Example of the show interfaces pfc summary Command

Example of the show interface pfc statistics Command

Example of the show interface ets summary Command

Example of the show interface ets detail Command

Example of the show stack-unit all stack-ports all pfc details Command

Example of the show stack-unit all stack-ports all ets details Command

Example of the show interface DCBx detail Command

```
Dell(conf)# show dot1p-queue-mapping
Dot1p Priority: 0 1 2 3 4 5 6 7
Queue          : 0 0 0 1 2 3 3 3
```

```
Dell# show dcb
stack-unit 0 port-set 0
    DCB Status : Enabled
    PFC Port Count : 56 (current), 56 (configured)
    PFC Queue Count : 2 (current), 2 (configured)
```

```
Dell# show interfaces tengigabitethernet 0/49 pfc summary
Interface TenGigabitEthernet 0/49
  Admin mode is on
  Admin is enabled
  Remote is enabled, Priority list is 4
  Remote Willing Status is enabled
  Local is enabled
  Oper status is Recommended
  PFC DCBx Oper status is Up
  State Machine Type is Feature
  TLV Tx Status is enabled
  PFC Link Delay 45556 pause quantams
  Application Priority TLV Parameters :
  -----
  FCOE TLV Tx Status is disabled
  ISCSI TLV Tx Status is disabled
  Local FCOE PriorityMap is 0x8
  Local ISCSI PriorityMap is 0x10
  Remote FCOE PriorityMap is 0x8
  Remote ISCSI PriorityMap is 0x8
```

```
Dell# show interfaces tengigabitethernet 0/49 pfc detail
Interface TenGigabitEthernet 0/49
  Admin mode is on
  Admin is enabled
  Remote is enabled
  Remote Willing Status is enabled
  Local is enabled
  Oper status is recommended
  PFC DCBx Oper status is Up
  State Machine Type is Feature
  TLV Tx Status is enabled
```

```
PFC Link Delay 45556 pause quanta
Application Priority TLV Parameters :
```

```
-----
FCOE TLV Tx Status is disabled
ISCSI TLV Tx Status is disabled
Local FCOE PriorityMap is 0x8
Local ISCSI PriorityMap is 0x10
Remote FCOE PriorityMap is 0x8
Remote ISCSI PriorityMap is 0x8
```

0 Input TLV pkts, 1 Output TLV pkts, 0 Error pkts, 0 Pause Tx pkts, 0 Pause Rx pkts

The following table describes the `show interface pfc summary` command fields.

Table 80. show interface pfc summary Command Description

Fields	Description
Interface	Interface type with stack-unit and port number.
Admin mode is on; Admin is enabled	PFC Admin mode is on or off with a list of the configured PFC priorities . When PFC admin mode is on, PFC advertisements are enabled to be sent and received from peers; received PFC configuration takes effect. The admin operational status for a DCBx exchange of PFC configuration is enabled or disabled.
Remote is enabled; Priority list Remote Willing Status is enabled	Operational status (enabled or disabled) of peer device for DCBx exchange of PFC configuration with a list of the configured PFC priorities. Willing status of peer device for DCBx exchange (Willing bit received in PFC TLV): enabled or disabled.
Local is enabled	DCBx operational status (enabled or disabled) with a list of the configured PFC priorities
Operational status (local port)	DCBx operational status (enabled or disabled) with a list of the configured PFC priorities. Port state for current operational PFC configuration: <ul style="list-style-type: none"> • Init: Local PFC configuration parameters were exchanged with peer. • Recommend: Remote PFC configuration parameters were received from peer. • Internally propagated: PFC configuration parameters were received from configuration source.
PFC DCBx Oper status	Operational status for exchange of PFC configuration on local port: match (up) or mismatch (down).
State Machine Type	Type of state machine used for DCBx exchanges of PFC parameters: <ul style="list-style-type: none"> • Feature: for legacy DCBx versions

Fields	Description
	<ul style="list-style-type: none"> Symmetric: for an IEEE version
TLV Tx Status	Status of PFC TLV advertisements: enabled or disabled.
PFC Link Delay	Link delay (in quanta) used to pause specified priority traffic.
Application Priority TLV: FCOE TLV Tx Status	Status of FCoE advertisements in application priority TLVs from local DCBx port: enabled or disabled.
Application Priority TLV: ISCSI TLV Tx Status	Status of ISCSI advertisements in application priority TLVs from local DCBx port: enabled or disabled.
Application Priority TLV: Local FCOE Priority Map	Priority bitmap used by local DCBx port in FCoE advertisements in application priority TLVs.
Application Priority TLV: Local ISCSI Priority Map	Priority bitmap used by local DCBx port in ISCSI advertisements in application priority TLVs.
Application Priority TLV: Remote FCOE Priority Map	Status of FCoE advertisements in application priority TLVs from remote peer port: enabled or disabled.
Application Priority TLV: Remote ISCSI Priority Map	Status of iSCSI advertisements in application priority TLVs from remote peer port: enabled or disabled.
PFC TLV Statistics: Input TLV pkts	Number of PFC TLVs received.
PFC TLV Statistics: Output TLV pkts	Number of PFC TLVs transmitted.
PFC TLV Statistics: Error pkts	Number of PFC error packets received.
PFC TLV Statistics: Pause Tx pkts	Number of PFC pause frames transmitted.
PFC TLV Statistics: Pause Rx pkts	Number of PFC pause frames received

```
Dell#show interfaces tengigabitethernet 0/3 pfc statistics
Interface TenGigabitEthernet 0/3
```

```
Priority Rx XOFF Frames Rx Total Frames Tx Total Frames
-----
```

0	0	0	0
1	0	0	0
2	0	0	0
3	0	0	0
4	0	0	0
5	0	0	0
6	0	0	0
7	0	0	0

```
Dell(conf)# show interfaces te 0/0 ets summary
Interface TenGigabitEthernet 0/0
Max Supported TC Groups is 4
```

Number of Traffic Classes is 8

Admin mode is on

Admin Parameters :

Admin is enabled

TC-grp	Priority#	Bandwidth	TSA
0	0,1,2,3,4,5,6,7	100%	ETS
1		0%	ETS
2		0%	ETS
3		0%	ETS
4		0%	ETS
5		0%	ETS
6		0%	ETS
7		0%	ETS

Priority#	Bandwidth	TSA
0	13%	ETS
1	13%	ETS
2	13%	ETS
3	13%	ETS
4	12%	ETS
5	12%	ETS
6	12%	ETS
7	12%	ETS

Remote Parameters:

Remote is disabled

Local Parameters :

Local is enabled

TC-grp	Priority#	Bandwidth	TSA
0	0,1,2,3,4,5,6,7	100%	ETS
1		0%	ETS
2		0%	ETS
3		0%	ETS
4		0%	ETS
5		0%	ETS
6		0%	ETS
7		0%	ETS

Priority#	Bandwidth	TSA
0	13%	ETS
1	13%	ETS
2	13%	ETS
3	13%	ETS
4	12%	ETS
5	12%	ETS
6	12%	ETS
7	12%	ETS

Oper status is init

Conf TLV Tx Status is disabled

Traffic Class TLV Tx Status is disabled

0 Input Conf TLV Pkts, 0 Output Conf TLV Pkts, 0 Error Conf TLV Pkts

0 Input Traffic Class TLV Pkts, 0 Output Traffic Class TLV Pkts, 0 Error

Traffic Class TLV

Pkts

The following table describes the show interface ets detail command fields.

```
Dell(conf)# show interfaces tengigabitethernet 0/0 ets detail
Interface TenGigabitEthernet 0/0
```

Max Supported TC Groups is 4
Number of Traffic Classes is 8
Admin mode is on
Admin Parameters :

Admin is enabled

TC-grp	Priority#	Bandwidth	TSA
0	0,1,2,3,4,5,6,7	100%	ETS
1		0%	ETS
2		0%	ETS
3		0%	ETS
4		0%	ETS
5		0%	ETS
6		0%	ETS
7		0%	ETS

Priority#	Bandwidth	TSA
0	13%	ETS
1	13%	ETS
2	13%	ETS
3	13%	ETS
4	12%	ETS
5	12%	ETS
6	12%	ETS
7	12%	ETS

Remote Parameters:

Remote is disabled

Local Parameters :

Local is enabled

TC-grp	Priority#	Bandwidth	TSA
0	0,1,2,3,4,5,6,7	100%	ETS
1		0%	ETS
2		0%	ETS
3		0%	ETS
4		0%	ETS
5		0%	ETS
6		0%	ETS
7		0%	ETS

Priority#	Bandwidth	TSA
0	13%	ETS
1	13%	ETS
2	13%	ETS
3	13%	ETS
4	12%	ETS
5	12%	ETS
6	12%	ETS
7	12%	ETS

Oper status is init

Conf TLV Tx Status is disabled

Traffic Class TLV Tx Status is disabled

0 Input Conf TLV Pkts, 0 Output Conf TLV Pkts, 0 Error Conf TLV Pkts

0 Input Traffic Class TLV Pkts, 0 Output Traffic Class TLV Pkts, 0 Error

Traffic Class TLV

Pkts

Table 81. show interface ets detail Command Description

Field	Description
Interface	Interface type with stack-unit and port number.
Max Supported TC Group	Maximum number of priority groups supported.
Number of Traffic Classes	Number of 802.1p priorities currently configured.
Admin mode	ETS mode: on or off. When on, the scheduling and bandwidth allocation configured in an ETS output policy or received in a DCBx TLV from a peer can take effect on an interface.
Admin Parameters	ETS configuration on local port, including priority groups, assigned dot1p priorities, and bandwidth allocation.
Remote Parameters	ETS configuration on remote peer port, including Admin mode (enabled if a valid TLV was received or disabled), priority groups, assigned dot1p priorities, and bandwidth allocation. If the ETS Admin mode is enabled on the remote port for DCBx exchange, the Willing bit received in ETS TLVs from the remote peer is included.
Local Parameters	ETS configuration on local port, including Admin mode (enabled when a valid TLV is received from a peer), priority groups, assigned dot1p priorities, and bandwidth allocation.
Operational status (local port)	Port state for current operational ETS configuration: <ul style="list-style-type: none"> • Init: Local ETS configuration parameters were exchanged with peer. • Recommend: Remote ETS configuration parameters were received from peer. • Internally propagated: ETS configuration parameters were received from configuration source.
ETS DCBx Oper status	Operational status of ETS configuration on local port: match or mismatch.
State Machine Type	Type of state machine used for DCBx exchanges of ETS parameters: <ul style="list-style-type: none"> • Feature: for legacy DCBx versions • Asymmetric: for an IEEE version
Conf TLV Tx Status	Status of ETS Configuration TLV advertisements: enabled or disabled.
ETS TLV Statistic: Input Conf TLV pkts	Number of ETS Configuration TLVs received.
ETS TLV Statistic: Output Conf TLV pkts	Number of ETS Configuration TLVs transmitted.

Field	Description
ETS TLV Statistic: Error Conf TLV pkts	Number of ETS Error Configuration TLVs received.

Dell(conf)# show stack-unit all stack-ports all pfc details

```
stack unit 0 stack-port all
  Admin mode is On
  Admin is enabled, Priority list is 4-5
  Local is enabled, Priority list is 4-5
  Link Delay 45556 pause quantum
  0 Pause Tx pkts, 0 Pause Rx pkts
```

```
stack unit 1 stack-port all
  Admin mode is On
  Admin is enabled, Priority list is 4-5
  Local is enabled, Priority list is 4-5
  Link Delay 45556 pause quantum
  0 Pause Tx pkts, 0 Pause Rx pkts
```

Dell(conf)# show stack-unit all stack-ports all ets details

```
Stack unit 0 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1
```

```
Admin mode is on
Admin Parameters:
```

```
-----
Admin is enabled
```

TC-grp	Priority#	Bandwidth	TSA
0	0,1,2,3,4,5,6,7	100%	ETS
1		-	-
2		-	-
3		-	-
4		-	-
5		-	-
6		-	-
7		-	-
8		-	-

```
Stack unit 1 stack port all
Max Supported TC Groups is 4
Number of Traffic Classes is 1
```

```
Admin mode is on
Admin Parameters:
```

```
-----
Admin is enabled
```

TC-grp	Priority#	Bandwidth	TSA
0	0,1,2,3,4,5,6,7	100%	ETS
1		-	-
2		-	-
3		-	-
4		-	-
5		-	-
6		-	-
7		-	-
8		-	-

```

Dell(conf)# show interface tengigabitethernet 0/49 dcbx detail
Dell#show interface te 0/49 dcbx detail

E-ETS Configuration TLV enabled           e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled         r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled         p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled  f-Application Priority for FCOE
disabled
I-Application priority for iSCSI enabled  i-Application Priority for iSCSI
disabled

```

```
-----
--
```

```

Interface TenGigabitEthernet 0/49
  Remote Mac Address 00:00:00:00:00:11
  Port Role is Auto-Upstream
  DCBX Operational Status is Enabled
  Is Configuration Source? TRUE

```

```

Local DCBX Compatibility mode is CEE
Local DCBX Configured mode is CEE
Peer Operating version is CEE
Local DCBX TLVs Transmitted: ErPfi

```

Local DCBX Status

```

-----
DCBX Operational Version is 0
DCBX Max Version Supported is 0
Sequence Number: 2
Acknowledgment Number: 2
Protocol State: In-Sync

```

Peer DCBX Status:

```

-----
DCBX Operational Version is 0
DCBX Max Version Supported is 255
Sequence Number: 2
Acknowledgment Number: 2
Total DCBX Frames transmitted 27
Total DCBX Frames received 6
Total DCBX Frame errors 0
Total DCBX Frames unrecognized 0

```

The following table describes the show interface DCBx detail command fields.

Table 82. show interface DCBx detail Command Description

Field	Description
Interface	Interface type with chassis slot and port number.
Port-Role	Configured DCBx port role: auto-upstream, auto-downstream, config-source, or manual.
DCBx Operational Status	Operational status (enabled or disabled) used to elect a configuration source and internally propagate a DCB configuration. The DCBx operational status is the combination of PFC and ETS operational status.

Field	Description
Configuration Source	Specifies whether the port serves as the DCBx configuration source on the switch: true (yes) or false (no).
Local DCBx Compatibility mode	DCBx version accepted in a DCB configuration as compatible. In auto-upstream mode, a port can only receive a DCBx version supported on the remote peer.
Local DCBx Configured mode	DCBx version configured on the port: CEE, CIN, IEEE v2.5, or Auto (port auto-configures to use the DCBx version received from a peer).
Peer Operating version	DCBx version that the peer uses to exchange DCB parameters.
Local DCBx TLVs Transmitted	Transmission status (enabled or disabled) of advertised DCB TLVs (see TLV code at the top of the show command output).
Local DCBx Status: DCBx Operational Version	DCBx version advertised in Control TLVs.
Local DCBx Status: DCBx Max Version Supported	Highest DCBx version supported in Control TLVs.
Local DCBx Status: Sequence Number	Sequence number transmitted in Control TLVs.
Local DCBx Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs.
Local DCBx Status: Protocol State	Current operational state of DCBx protocol: ACK or IN-SYNC.
Peer DCBx Status: DCBx Operational Version	DCBx version advertised in Control TLVs received from peer device.
Peer DCBx Status: DCBx Max Version Supported	Highest DCBx version supported in Control TLVs received from peer device.
Peer DCBx Status: Sequence Number	Sequence number transmitted in Control TLVs received from peer device.
Peer DCBx Status: Acknowledgment Number	Acknowledgement number transmitted in Control TLVs received from peer device.
Total DCBx Frames transmitted	Number of DCBx frames sent from local port.
Total DCBx Frames received	Number of DCBx frames received from remote peer port.
Total DCBx Frame errors	Number of DCBx frames with errors received.
Total DCBx Frames unrecognized	Number of unrecognizable DCBx frames received.

PFC and ETS Configuration Examples

This section contains examples of how to configure and apply DCB input and output policies on an interface.

Using PFC and ETS to Manage Data Center Traffic

The following shows examples of using PFC and ETS to manage your data center traffic.

In the following example:

- Incoming SAN traffic is configured for priority-based flow control.
- Outbound LAN, IPC, and SAN traffic is mapped into three ETS priority groups and configured for enhanced traffic selection (bandwidth allocation and scheduling).
- One lossless queue is used.

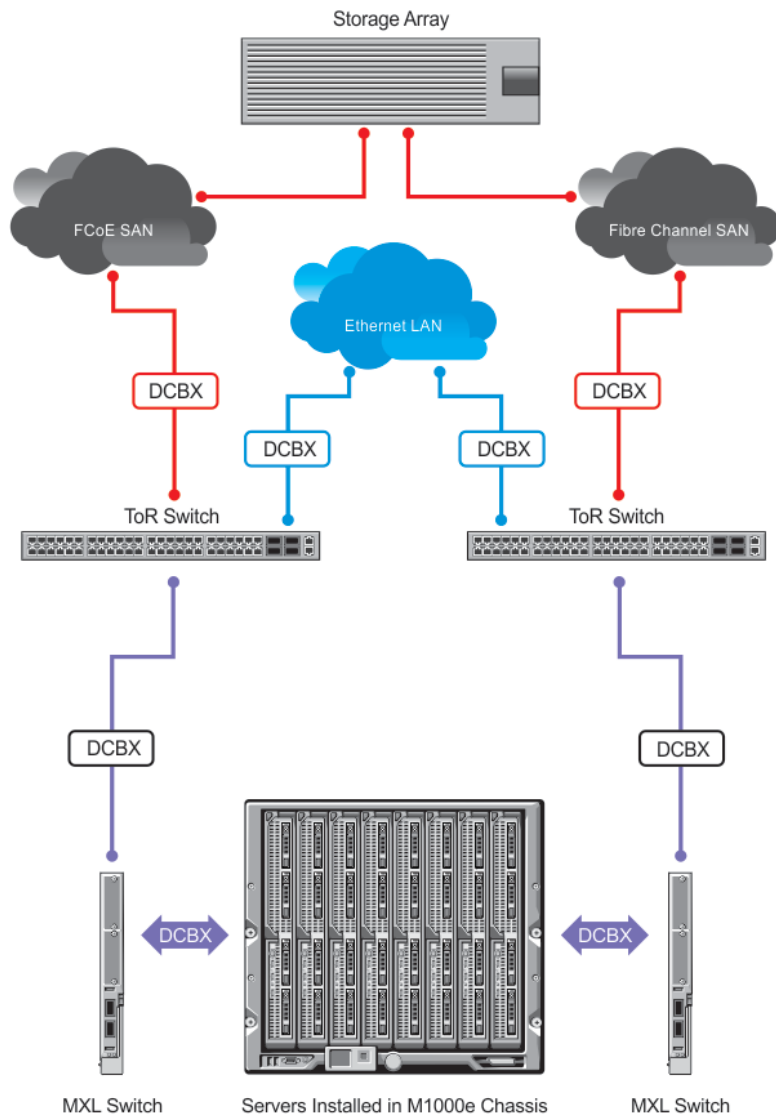


Figure 143. PFC and ETS Applied to LAN, IPC, and SAN Priority Traffic

QoS Traffic Classification: The `service-class dynamic dot1p` command has been used in Global Configuration mode to map ingress dot1p frames to the queues shown in the following table. For more information, refer to [QoS dot1p Traffic Classification and Queue Assignment](#).

dot1p Value in Incoming Frame	Queue Assignment
0	0
1	0
2	0
3	1

dot1p Value in Incoming Frame	Queue Assignment
4	2
5	3
6	3
7	3

The following describes the dot1p-priority class group assignment

dot1p Value in the Incoming Frame	Priority Group Assignment
0	LAN
1	LAN
2	LAN
3	SAN
4	IPC
5	LAN
6	LAN
7	LAN

The following describes the priority group-bandwidth assignment.

Priority Group	Bandwidth Assignment
IPC	5%
SAN	50%
LAN	45%

PFC and ETS Configuration Command Examples

The following examples show PFC and ETS configuration commands to manage your data center traffic.

Example of Configuring QoS Priority-Queue Assignment to Honor Dot1p Priorities

```
Dell(conf)# service-class dynamic dot1p
```

Using PFC and ETS to Manage Converged Ethernet Traffic in a Switch Stack

The following example shows how to apply the DCB PFC input policy (`ipc_san_lan`) and ETS output policy (`ets`) on all MXL switches in a switch stack.

This example references the [PFC and ETS Configuration Examples](#) section.

Example of Applying DCB PFC Input Policy and ETS Output Policy in a Switch Stack

```
dcb-map stack-unit all stack-ports all <dcb-map-name>
```

Interworking of DCB Map With DCB Buffer Threshold Settings

The `dcb-input` and `dcb-output` configuration commands are deprecated. You must use the `dcb-map` command to create a DCB map to configure priority flow control (PFC) and enhanced transmission selection (ETS) on Ethernet ports that support converged Ethernet traffic.

Configure the `dcb-buffer-threshold` command and its related parameters only on ports with either auto configuration or `dcb-map` configuration. This command is not supported on existing front-panel interfaces or stack ports that are configured with the `dcb-input` or `dcb-output` commands. Similarly, if the `dcb-buffer-threshold` configuration is present on a stack port or any interface, the `dcb-input` or `dcb-output` policies cannot be applied on those interfaces.

Example: When the `dcb-buffer-threshold` policy is applied on interfaces or stack ports with the `dcb-input` or `dcb-output` policies, the following error message is displayed:

```
%Error: dcb-buffer-threshold not supported on interfaces with deprecated
commands
```

Example: When the `dcb-input` or `dcb-output` policy is configured on interfaces or stack ports with the `dcb-buffer` threshold policy, the following error message is displayed:

```
%Error: Deprecated command is not supported on interfaces with dcb-buffer-
threshold configured
```

You must not modify the `service-class dot1p` mappings when any `buffer-threshold-policy` is configured on the system.

```
Dell(conf)#service-class dot1p-mapping dot1p0 3
```

```
% Error: PFC buffer-threshold policies conflict with dot1p mappings. Please
remove all dcb-buffer-threshold policies to change mappings.
```

The `show dcb` command has been enhanced to display the following additional buffer-related information:

```
Dell(conf)#do show dcb
dcb Status : Enabled
PFC Queue Count : 2 --Indicate the PFC queue configured.
Total buffer (lossy + lossless) (in KB): 7787--Total buffer space for lossy and
lossless queues
PFC total buffer (in KB): 6526 --Indicates the total buffer (configured or
default)
PFC shared buffer (in KB): 832--Indicates the shared buffer (Configured or
default)
PFC available buffer ( in KB): 5694--Indicates remaining available buffers for
PFC that are free to be allocated
```


Fibre Channel over Ethernet for FC Flex IO Modules

FCoE provides a converged Ethernet network that allows the combination of storage-area network (SAN) and LAN traffic on a Layer 2 link by encapsulating Fibre Channel data into Ethernet frames.

The Fibre Channel (FC) Flex IO module is supported on Dell Networking Operating System (OS) MXL 10/40GbE Switch and M I/O Aggregator (IOA). The MXL and IOA switch installed with the FC Flex IO module functions as a top-of-rack edge switch that supports converged enhanced Ethernet (CEE) traffic — Fibre channel over Ethernet (FCoE) for storage, Interprocess Communication (IPC) for servers, and Ethernet local area network (LAN) (IP cloud) for data — as well as FC links to one or more storage area network (SAN) fabrics.

FCoE works with the Ethernet enhancements provided in Data Center Bridging (DCB) to support lossless (no-drop) SAN and LAN traffic. In addition, DCB provides flexible bandwidth sharing for different traffic types, such as LAN and SAN, according to 802.1p priority classes of service. DCBx should be enabled on the system before the FIP snooping feature is enabled.

All of the commands that are supported for FCoE on the MXL and I/O Aggregator apply to the FC Flex IO modules. Similarly, all of the configuration procedures and the settings that are applicable for FCoE on the MXL and I/O Aggregator are valid for the FC Flex IO modules.

NPIV Proxy Gateway for FC Flex IO Modules

The N-port identifier virtualization (NPIV) Proxy Gateway (NPG) feature provides FCoE-FC bridging capability on the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module switch, allowing server CNAs to communicate with SAN fabrics over the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module.

To configure the M I/O Aggregator and MXL 10/40GbE Switch with the FC Flex IO module to operate as an NPIV proxy gateway, use the following commands:

NPIV Proxy Gateway Configuration on FC Flex IO Modules

The Fibre Channel (FC) Flex IO module is supported on the MXL 10/40GbE Switch and M I/O Aggregator (IOA). The MXL and IOA switches, installed with the FC Flex IO module, function as a top-of-rack edge switch that supports Converged Enhanced Ethernet (CEE) traffic — Fibre Channel over Ethernet (FCoE) for storage, Interprocess Communication (IPC) for servers, and Ethernet local area network (LAN) (IP cloud) for data — as well as FC links to one or more storage area network (SAN) fabrics.

The N-port identifier virtualization (NPIV) proxy gateway (NPG) provides FCoE-FC bridging capability on the MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module.

This chapter describes how to configure and use an NPIV proxy gateway on an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module in a (SAN).

NPIV Proxy Gateway Operations and Capabilities

Benefits of an NPIV Proxy Gateway

The MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module functions as a top-of-rack edge switch that supports Converged Enhanced Ethernet (CEE) traffic — FCoE for storage, Interprocess Communication (IPC) for servers, and Ethernet LAN (IP cloud) for data — as well as Fibre Channel (FC) links to one or more SAN fabrics.

Using an NPIV proxy gateway (NPG) helps resolve the following problems in a storage area network:

- Fibre Channel storage networks typically consist of servers connected to edge switches, which are connected to SAN core switches. As the SAN grows, it is necessary to add more ports and SAN switches. This results in an increase in the required domain IDs, which may surpass the upper limit of 239 domain IDs supported in the SAN network. An NPG avoids the need for additional domain IDs because it is deployed outside the SAN and uses the domain IDs of core switches in its FCoE links.
- With the introduction of 10GbE links, FCoE is being implemented for server connections to optimize performance. However, a SAN traditionally uses Fibre Channel to transmit storage traffic. FCoE servers require an efficient and scalable bridging feature to access FC storage arrays, which an NPG provides.

NPIV Proxy Gateway Operation

Consider a sample scenario of NPG operation. An M1000e chassis configured as an NPG does not join a SAN fabric, but functions as an FCoE-FC bridge that forwards storage traffic between servers and core SAN switches. The core switches forward SAN traffic to and from FC storage arrays.

An M1000e chassis FC port is configured as an N (node) port that logs in to an F (fabric) port on the upstream FC core switch and creates a channel for N-port identifier virtualization. NPIV allows multiple N-port fabric logins at the same time on a single, physical Fibre Channel link.

Converged Network Adapter (CNA) ports on servers connect to the M1000e chassis Ten-Gigabit Ethernet ports and log in to an upstream FC core switch through the FC Flex IO module N port. Server fabric login (FLOGI) requests are converted into fabric discovery (FDISC) requests before being forwarded by the FC Flex IO module to the FC core switch.

Servers use CNA ports to connect over FCoE to an Ethernet port in ENode mode on the NPIV proxy gateway. FCoE transit with FIP snooping is automatically enabled and configured on the M1000e gateway to prevent unauthorized access and data transmission to the SAN network (see FCoE Transit). FIP is used by server CNAs to discover an FCoE switch operating as an FCoE forwarder (FCF).

The NPIV proxy gateway aggregates multiple locally connected server CNA ports into one or more upstream N port links, conserving the number of ports required on an upstream FC core switch while providing an FCoE-to-FC bridging functionality. The upstream N ports on an M1000e can connect to the same or multiple fabrics.

Using an FCoE map applied to downstream (server-facing) Ethernet ports and upstream (fabric-facing) FC ports, you can configure the association between a SAN fabric and the FCoE VLAN that connects

servers over the NPIV proxy gateway to FC switches in the fabric. An FCoE map virtualizes the upstream SAN fabric as an FCF to downstream CNA ports on FCoE-enabled servers as follows:

- As soon as an FC N port comes online (no `shutdown` command), the NPG starts sending FIP multicast advertisements, which contain the fabric name derived from the 64-bit worldwide name (WWN) of the principal SAN switch. (The principal switch in a fabric is the FC switch with the lowest domain ID.)
- When you apply the FCoE map to a server-facing Ethernet port in ENode mode, ACLs are automatically configured to allow only FCoE traffic from servers that perform a successful FLOGI on the FC switch. All other traffic on the VLAN is denied.

You can specify one or more upstream N ports in an FCoE map. The FCoE map also contains the VLAN ID of the dedicated VLAN used to transmit FCoE traffic between the SAN fabric and servers.

NPIV Proxy Gateway: Protocol Services

An MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPG provides the following protocol services:

- Fibre Channel service to create N ports and log in to an upstream FC switch.
- FCoE service to perform:
- Virtualization of FC N ports on an NPG so that they appear as FCoE FCFs to downstream servers.
- NPIV service to perform the association and aggregation of FCoE servers to upstream F ports on core switches (through N ports on the NPG). Conversion of server FLOGIs and FDISCs, which are received over MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module ENode ports, are converted into FDISCs addressed to the upstream F ports on core switches.

NPIV Proxy Gateway Functionality

An MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPG provides the following functionality in a storage area network:

- FIP Snooping bridge that provides security for FCoE traffic using ACLs (see FCoE Transit chapter).
- FCoE gateway that provides FCoE-to-FC bridging. N-port virtualization using FCoE maps exposes upstream F ports as FCF ports to downstream server-facing ENode ports on the NPG (see FCoE Maps).

NPIV Proxy Gateway: Terms and Definitions

The following table describes the terms used in an NPG configuration on the MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module.

Table 83. MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPIV Proxy Gateway: Terms and Definitions

Term	Description
FC port	Fibre Channel port on an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module FC module that operates in autosensing, 2, 4,

Term	Description
	or 8-Gigabit mode. On an NPIV proxy gateway, an FC port can be used as a downlink for a server connection and an uplink for a fabric connection.
F port	Port mode of an FC port connected to an end node (N) port on an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPIV proxy gateway.
N port	Port mode of an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module FC port that connects to an F port on an FC switch in a SAN fabric. On an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPIV proxy gateway, an N port also functions as a proxy for multiple server CNA-port connections
ENode port	Port mode of a server-facing MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module Ethernet port that provides access to FCF functionality on a fabric.
CNA port	N-port functionality on an FCoE-enabled server port. A converged network adapter (CNA) can use one or more Ethernet ports. CNAs can encapsulate Fibre Channel frames in Ethernet for FCoE transport and de-encapsulate Fibre Channel frames from FCoE to native Fibre Channel.
DCB map	Template used to configure DCB parameters, including priority-based flow control (PFC) and enhanced transmission selection (ETS), on CEE ports.
Fibre Channel fabric	Network of Fibre Channel devices and storage arrays that interoperate and communicate.
FCF	Fibre Channel forwarder: FCoE-enabled switch that can forward FC traffic to both downstream FCoE and upstream FC devices. An NPIV proxy gateway functions as an FCF to export upstream F port configurations to downstream server CNA ports.
FC-MAP	FCoE MAC-address prefix — The unique 24-bit MAC address prefix in FCoE packets used to generate a fabric-provided MAC address (FPMA). The FPMA is required to send FCoE packets from a server to a SAN fabric.
FCoE map	Template used to configure FCoE and FC parameters on Ethernet and FC ports in a converged fabric.
FCoE VLAN	VLAN dedicated to carrying only FCoE traffic between server CNA ports and a SAN fabric. (FCoE traffic must travel in a VLAN.) When you apply an FCoE map on a port, FCoE is enabled on the port. All non-FCoE traffic is dropped on an FCoE VLAN.
FIP	FCoE Initialization Protocol: Layer 2 protocol for endpoint discovery, fabric login, and fabric association. FIP is used by server CNAs to discover an upstream FCoE switch operating as an FCF. FIP keepalive messages maintain the connection between an FCoE initiator and an FCF.
NPIV	N-port identifier virtualization: The capability to map multiple FCoE links from downstream ports to a single upstream FC link.

Term	Description
principal switch	The switch in a fabric with the lowest domain number. The principal switch accesses the master name database and the zone/zone set database.

DCB Maps

A Data Center Bridging (DCB) map is used to configure DCB functionality, such as PFC and ETS, on MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module Ethernet ports that support CEE traffic and are DCBx-enabled, by default. For more information, on PFC and ETS, see Data Center Bridging (DCB).

By default, no PFC and ETS settings in a DCB map are applied to MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module Ethernet ports when they are enabled. On an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPG, you must configure PFC and ETS parameters in a DCB map and then apply the map to server-facing Ethernet ports (see the “Creating a DCB map” section).

FCoE Maps

An FCoE map is used to identify the SAN fabric to which FCoE storage traffic is sent. Using an FCoE map, an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module NPG operates as an FCoE-FC bridge between an FC SAN and FCoE network by providing FCoE-enabled servers and switches with the necessary parameters to log in to a SAN fabric.

An FCoE map applies the following parameters on server-facing Ethernet and fabric-facing FC ports on the MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module:

- The dedicated FCoE VLAN used to transport FCoE storage traffic.
- The FC-MAP value used to generate a fabric-provided MAC address.
- The association between the FCoE VLAN ID and FC fabric ID where the desired storage arrays are installed. Each Fibre Channel fabric serves as an isolated SAN topology within the same physical network.
- The priority used by a server to select an upstream FCoE forwarder (FCF priority)
- FIP keepalive (FKA) advertisement timeout



NOTE:

In each FCoE map, the fabric ID, FC-MAP value, and FCoE VLAN must be unique. Use one FCoE map to access one SAN fabric. You cannot use the same FCoE map to access different fabrics.

When you configure an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module as an NPG, FCoE transit with FIP snooping is automatically enabled and configured using the parameters in the FCoE map applied to server-facing Ethernet and fabric-facing FC interfaces (see FIP Snooping on an NPIV Proxy Gateway).

After you apply an FCoE map on an FC port, when you enable the port (no shutdown), the NPG starts sending FIP multicast advertisements on behalf of the FC port to downstream servers in order to advertise the availability of a new FCF port on the FCoE VLAN. The FIP advertisement also contains a keepalive message to maintain connectivity between a SAN fabric and downstream servers.

Configuring an NPIV Proxy Gateway

Prerequisite: Before you configure an NPIV proxy gateway (NPG) with the FC Flex IO module on an MXL 10/40GbE Switch or an M I/O Aggregator, ensure that the following features are enabled.

- DCB is enabled by default with the FC Flex IO module on the MXL 10/40GbE Switch or M I/O Aggregator.
- Autonegotiated DCBx is enabled for converged traffic by default with the FC Flex IO module Ethernet ports on all MXL 10/40GbE Switches or M I/O Aggregators.
- FCoE transit with FIP snooping is automatically enabled when you configure Fibre Channel with the FC Flex IO module on the MXL 10/40GbE Switch or M I/O Aggregator.

To configure an NPG operation with the FC Flex IO module on an MXL 10/40GbE Switch or an M I/O Aggregator, follow these general configuration steps:

1. Enabling Fibre Channel Capability on the Switch
2. Creating a DCB map
3. Applying a DCB map on server-facing Ethernet ports
4. Creating an FCoE VLAN
5. Creating an FCoE map
6. Applying an FCoE map on server-facing Ethernet ports
7. Applying an FCoE Map on fabric-facing FC ports

Enabling Fibre Channel Capability on the Switch

Enable the FC Flex IO module on an MXL 10/40GbE Switch or an M I/O Aggregator that you want to configure as an NPG for the Fibre Channel protocol. When you enable Fibre Channel capability, FCoE transit with FIP snooping is automatically enabled on all VLANs on the switch, using the default FCoE transit settings.

Task	Command	Command Mode
Enable an MXL 10/40GbE Switch and M I/O Aggregator with the FC Flex IO module for the Fibre Channel protocol.	<code>feature fc</code>	CONFIGURATION

Creating a DCB Map

Configure the priority-based flow control (PFC) and enhanced traffic selection (ETS) settings in a DCB map before you apply them on downstream server-facing ports on an MXL 10/40GbE Switch or an M I/O Aggregator with the FC Flex IO module.

Step	Task	Command	Command Mode
1	Create a DCB map to specify PFC and ETS settings for groups of dot1p priorities.	<code>dcb-map name</code>	CONFIGURATION
2	<p>Configure the PFC setting (on or off) and the ETS bandwidth percentage allocated to traffic in each priority group. Configure whether the priority group traffic should be handled with strict-priority scheduling. The sum of all allocated bandwidth percentages must be 100 percent. Strict-priority traffic is serviced first. Afterward, bandwidth allocated to other priority groups is made available and allocated according to the specified percentages. If a priority group does not use its allocated bandwidth, the unused bandwidth is made available to other priority groups.</p> <p>Restriction: You can enable PFC on a maximum of two priority queues.</p> <p>Repeat this step to configure PFC and ETS traffic handling for each priority group, for example:</p> <pre>priority-group 0 bandwidth 60 pfc off priority-group 1 bandwidth 20 pfc on priority-group 2 bandwidth 20 pfc on priority-group 4 strict-priority pfc off</pre>	<pre>priority-group group_num {bandwidth percentage strict- priority} pfc {on off}</pre>	DCB MAP
3	<p>Specify the priority group ID number to handle VLAN traffic for each dot1p class-of-service: 0 through 7. Leave a space between each priority group number. For example, <code>priority-pgid 0 0 0 1 2 4 4 4</code> where dot1p priorities 0, 1, and 2 are mapped to priority group 0; dot1p priority 3 is mapped to priority group 1; dot1p priority 4 is mapped to priority group 2; dot1p priorities 5, 6, and 7 are mapped to priority group 4.</p> <p>All priorities that map to the same egress queue must be in the same priority group.</p>	<pre>priority-pgid dot1p0_group_num dot1p1_group_num dot1p2_group_num dot1p3_group_num dot1p4_group_num dot1p5_group_num dot1p6_group_num dot1p7_group_num</pre>	DCB MAP

Important Points to Remember

- If you remove a dot1p priority-to-priority group mapping from a DCB map (no `priority pgid` command), the PFC and ETS parameters revert to their default values on the interfaces on which the DCB map is applied. By default, PFC is not applied on specific 802.1p priorities; ETS assigns equal bandwidth to each 802.1p priority.
As a result, PFC and lossless port queues are disabled on 802.1p priorities, and all priorities are mapped to the same priority queue and equally share port bandwidth.
- To change the ETS bandwidth allocation configured for a priority group in a DCB map, do not modify the existing DCB map configuration. Instead, create a new DCB map with the desired PFC and ETS settings, and apply the new map to the interfaces to override the previous DCB map settings. Then, delete the original dot1p priority-to-priority group mapping.

If you delete the dot1p priority-to-priority group mapping (`no priority ppid` command) before you apply the new DCB map, the default PFC and ETS parameters are applied on the interfaces. This change may create a DCB mismatch with peer DCB devices and interrupt the network operation.

Applying a DCB Map on Server-facing Ethernet Ports

You can apply a DCB map only on a physical Ethernet interface and can apply only one DCB map per interface.

Step	Task	Command	Command Mode
1	<p>Enter CONFIGURATION mode on a server-facing port or port channel to apply a DCB map.</p> <p>You cannot apply a DCB map on a port channel. However, you can apply a DCB map on the ports that are members of the port channel.</p>	<pre>interface {tengigabitEthernet slot/port fortygigabitEthernet slot/port}</pre>	CONFIGURATION
2	<p>Apply the DCB map on an Ethernet port or port channel. The port is configured with the PFC and ETS settings in the DCB map, for example:</p> <pre>Dell# interface tengigabitEthernet 0/0 Dell(config-if-te-0/0)# dcb-map SAN_DCB1</pre> <p>Repeat this step to apply a DCB map to more than one port or port channel.</p>	<pre>dcb-map name</pre>	INTERFACE

Creating an FCoE VLAN

Create a dedicated VLAN to send and receive Fibre Channel traffic over FCoE links between servers and a fabric over an NPG. The NPG receives FCoE traffic and forwards decapsulated FC frames over FC links to SAN switches in a specified fabric.

Step	Task	Command	Command Mode
1	<p>Create the dedicated VLAN for FCoE traffic.</p> <p>Range: 2-4094.</p> <p>VLAN 1002 is commonly used to transmit FCoE traffic.</p>	<pre>interface vlan vlan- id</pre>	CONFIGURATION

When you apply an FCoE map to an Ethernet port (Applying an FCoE map on server-facing Ethernet ports), the port is automatically configured as a tagged member of the FCoE VLAN.

Creating an FCoE Map

An FCoE map consists of:

- An association between the dedicated VLAN, used to carry FCoE traffic, and the SAN fabric where the storage arrays are installed. Use a separate FCoE VLAN for each fabric to which the FCoE traffic is forwarded. Any non-FCoE traffic sent on a dedicated FCoE VLAN is dropped.
- The FC-MAP value, used to generate the fabric-provided MAC address (FPMA). The FPMA is used by servers to transmit FCoE traffic to the fabric. You can associate an FC-MAP with only one FCoE VLAN and conversely, associate an FCoE VLAN with only one FC-MAP.
- FCF priority, the priority used by a server CNA to select an upstream FCoE forwarder (FCF)
- FIP keepalive (FKA) advertisement timeout

The values for the FCoE VLAN, fabric ID and FC-MAP must be unique. Apply an FCoE map on downstream server-facing Ethernet ports and upstream fabric-facing Fibre Channel ports.

Step	Task	Command	Command Mode
1	Create an FCoE map that contains parameters used in the communication between servers and a SAN fabric.	<code>fcoe-map map-name</code>	CONFIGURATION
2	Configure the association between the dedicated VLAN and the fabric where the desired storage arrays are installed. The fabric and VLAN ID numbers must be the same. Fabric and VLAN ID range: 2-4094. For example: <code>fabric id 10 vlan 10</code>	<code>fabric-id fabric-num</code> <code>vlan vlan-id</code>	FCoE MAP
3	Add a text description of the settings in the FCoE map. Maximum: 32 characters.	<code>description text</code>	FCoE MAP
4	Specify the FC-MAP value used to generate a fabric-provided MAC address, which is required to send FCoE traffic from a server on the FCoE VLAN to the FC fabric specified in Step 2. Enter a unique MAC address prefix as the FC-MAP value for each fabric. Range: 0EFC00–0EFCFF. Default: None.	<code>fc-map fc-map-value</code>	FCoE MAP
5	Configure the priority used by a server CNA to select the FCF for a fabric login (FLOGI). Range: 1-255. Default: 128.	<code>fcf-priority</code> <code>priority</code>	FCoE MAP
6	Enable the monitoring FIP keep-alive messages (if it is disabled) to detect if other	<code>keepalive</code>	FCoE MAP

FCoE devices are reachable. Default: FIP keep-alive monitoring is enabled.

- | | | | |
|---|---|---|----------|
| 7 | Configure the time interval (in seconds) used to transmit FIP keepalive advertisements. | <code>fka-adv-period</code>
<code>seconds</code> | FCoE MAP |
|---|---|---|----------|

Range: 8-90 seconds. Default: 8 seconds.

Applying an FCoE Map on Server-facing Ethernet Ports

You can apply multiple FCoE maps on an Ethernet port or port channel. When you apply an FCoE map on a server-facing port or port channel:

- The port is configured to operate in hybrid mode (accept both tagged and untagged VLAN frames).
- The associated FCoE VLAN is enabled on the port or port channel.

When you enable a server-facing Ethernet port, the servers respond to the FIP advertisements by performing FLOGIs on upstream virtualized FCF ports. The NPG forwards the FLOGIs as fabric discovery (FDISC) messages to a SAN switch.

Step	Task	Command	Command Mode
1	Configure a server-facing Ethernet port or port channel with an FCoE map.	<code>interface</code> { tengigabitEthernet <code>slot/port</code> fortygigabitEthernet <code>slot/port</code> port- channel <code>num</code> }	CONFIGURATION
2	Apply the FCoE/FC configuration in an FCoE map on the Ethernet port. Repeat this step to apply an FCoE map to more than one port, for example:	<code>fcoe-map map-name</code>	INTERFACE or INTERFACE PORT_CHANNEL
	<pre>Dell# interface tengigabitEthernet 0/0Dell(config-if-te-0/0)# fcoe-map SAN_FABRIC_A Dell# interface port- channel 3Dell(config-if-po-3)# fcoe-map SAN_FABRIC_A Dell# interface fortygigabitEthernet 0/48Dell(config-if-fo-0/0)# fcoe- map SAN_FABRIC_A</pre>		
3	Enable the port for FCoE transmission using the map settings.	<code>no shutdown</code>	INTERFACE

Applying an FCoE Map on Fabric-facing FC Ports

The MXL 10/40GbE Switch and M I/O Aggregator, with the FC Flex IO module FC ports, are configured by default to operate in N port mode to connect to an F port on an FC switch in a fabric. You can apply only one FCoE map on an FC port.

When you apply an FCoE map on a fabric-facing FC port, the FC port becomes part of the FCoE fabric, whose settings in the FCoE map are configured on the port and exported to downstream server CNA ports.

Each MXL 10/40GbE Switch and M I/O Aggregator, with the FC Flex IO module FC port, is associated with an Ethernet MAC address (FCF MAC address). When you enable a fabric-facing FC port, the FCoE map applied to the port starts sending FIP multicast advertisements using the parameters in the FCoE map over server-facing Ethernet ports. A server sees the FC port, with its applied FCoE map, as an FCF port.

Step	Task	Command	Command Mode
1	Configure a fabric-facing FC port.	<code>interface fibrechannel slot/ port</code>	CONFIGURATION
2	Apply the FCoE and FC fabric configurations in an FCoE map to the port. Repeat this step to apply an FCoE map to more than one FC port, for example: Dell# interface fi 0/0Dell(config-if-fc-0/0)# fabric SAN_FABRIC_A	<code>fabric map-name</code>	INTERFACE FIBRE_CHANNEL
3	Enable the port for FC transmission.	<code>no shutdown</code>	INTERFACE FIBRE_CHANNEL

Tip: You can apply a DCB or FCoE map to a range of Ethernet or Fibre Channel interfaces by using the `interface range` command; for example:

```
Dell(config)# interface range tengigabitEthernet 1/12 - 23 , tengigabitEthernet 2/24 - 35
```

```
Dell(config)# interface range fibrechannel 0/0 - 3 , fibrechannel 0/8 - 11
```

Enter the keywords `interface range` followed by an interface type and port range. A port range must contain spaces before and after the dash. Separate each interface type and port range with a space, comma, and space as shown in the preceding examples.

Sample Configuration

1. Configure a DCB map with PFC and ETS settings:

```
Dell(config)# dcb-map SAN_DCB_MAP
Dell(config-dcbx-name)# priority-group 0 bandwidth 60 pfc off

Dell(config-dcbx-name)# priority-group 1 bandwidth 20 pfc on

Dell(config-dcbx-name)# priority-group 2 bandwidth 20 pfc on

Dell(config-dcbx-name)# priority-group 4 strict-priority pfc off

Dell(conf-dcbx-name)# priority-pgid 0 0 0 1 2 4 4 4
```

2. Apply the DCB map on a downstream (server-facing) Ethernet port:

- ```
Dell(config)# interface tengigabitethernet 1/0
```
- ```
Dell(config-if-te-0/0)#dcb-map SAN_DCB_MAP
```
3. Create the dedicated VLAN to be used for FCoE traffic:


```
Dell(conf)#interface vlan 1002
```
 4. Configure an FCoE map to be applied on downstream (server-facing) Ethernet and upstream (core-facing) FC ports:


```
Dell(config)# fcoe-map SAN_FABRIC_A
```

```
Dell(config-fcoe-name)# fabric-id 1002 vlan 1002
```

```
Dell(config-fcoe-name)# description "SAN_FABRIC_A"
```

```
Dell(config-fcoe-name)# fc-map 0efc00
```

```
Dell(config-fcoe-name)# keepalive
```

```
Dell(config-fcoe-name)# fcf-priority 128
```

```
Dell(config-fcoe-name)# fka-adv-period 8
```
 5. Enable an upstream FC port:


```
Dell(config)# interface fibrechannel 0/0
```

```
Dell(config-if-fc-0)# no shutdown
```
 6. Enable a downstream Ethernet port:


```
Dell(config)#interface tengigabitEthernet 0/0
```


```
Dell(conf-if-te-0)# no shutdown
```

Displaying NPIV Proxy Gateway Information

To display information on the NPG operation, use the show commands in the following table

Table 84. Displaying NPIV Proxy Gateway Information

Command	Description
<code>show interfaces status</code>	Displays the operational status of Ethernet and Fibre Channel interfaces on an MXL 10/40GbE Switch or M I/O Aggregator with the FC Flex IO module, NPG.

Command	Description
	 NOTE: Although the <code>show interface status</code> command displays the Fiber Channel (FC) interfaces with the abbreviated label of 'Fc' in the output, if you attempt to specify a FC interface by using the <code>interface fc</code> command in the CLI interface, an error message is displayed. You must configure FC interfaces by using the <code>interface fi</code> command in CONFIGURATION mode.
<code>show fcoe-map [brief map-name]</code>	<p>Displays the Fibre Channel and FCoE configuration parameters in FCoE maps. Enter the brief keyword to display an overview of currently configured FCoE maps.</p> <p>Enter the name of an FCoE map to display the FC and FCoE parameters configured in the map to be applied on MXL 10/40GbE Switch or M I/O Aggregator with the FC Flex IO module Ethernet (FCoE) and FC ports.</p>
<code>show qos dcb-map map-name</code>	Displays configuration parameters in a specified DCB map.
<code>show npiv devices [brief]</code>	Displays information on FCoE and FC devices currently logged in to the NPG.
<code>show fc switch</code>	Displays the FC mode of operation and worldwide node (WWN) name of an MXL 10/40GbE Switch or M I/O Aggregator with the FC Flex IO module.

show interfaces status Command Example

```
Dell# show interfaces status
Port      Description  Status Speed      Duplex  Vlan
Fc 0/0    Up          8000 Mbit Auto     --
Fc 0/1    Down       Auto      Auto     --
Fc 0/2    Down       Auto      Auto     --
Fc 0/3    Down       Auto      Auto     --
Fc 0/4    Down       Auto      Auto     --
Fc 0/5    Down       Auto      Auto     --
Fc 0/6    Down       Auto      Auto     --
Fc 0/7    Down       Auto      Auto     --
Fc 0/8    Down       Auto      Auto     --
Fc 0/9    Down       Auto      Auto     --
Fc 0/10   Down       Auto      Auto     --
Fc 0/11   Down       Auto      Auto     --
Te 1/12   Down       Auto      Auto     --
Te 1/13   Down       Auto      Auto     --
Te 1/14   Down       Auto      Auto     --
Te 1/15   Down       Auto      Auto     --
Te 1/16   Down       Auto      Auto     --
Te 1/17   Down       Auto      Auto     --
Te 1/18   Down       Auto      Auto     --
Te 1/19   Up          10000 Mbit Full  --
Te 1/20   Down       Auto      Auto     --
Te 1/21   Down       Auto      Auto     --
```

Table 85. show interfaces status Field Descriptions

Field	Description
Port	Server-facing 10GbE Ethernet (Te), 40GbE Ethernet (Fo), or fabric-facing Fibre Channel (Fc) port with <i>slot/port</i> information.
Description	Text description of port.
Status	Operational status of port: Ethernet ports - up (transmitting FCoE and LAN storage traffic) or down (not transmitting traffic). Fibre Channel ports - up (link is up and transmitting FC traffic) or down (link is down and not transmitting FC traffic), link-wait (link is up and waiting for FLOGI to complete on peer SW port), or removed (port has been shut down).
Speed	Transmission speed (in Megabits per second) of Ethernet and FC iports, including auto-negotiated speed (Auto).
Duplex	Data transmission mode: Full (allows communication in both directions at the same time), Half (allows communication in both directions but not at the same time), Auto (auto-negotiated transmission).
VLAN	VLAN IDs of the VLANs in which the port is a member.

show fcoe-map Command Examples

```
Dell# show fcoe-map brief
Fabric-Name  Fabric-Id      Vlan-Id  FC-MAP    FCF-Priority  Config-
State      Oper-State
fid_1003    1003                1003     0efc03    128
ACTIVE      UP
fid_1004    1004                1004     0efc04    128
ACTIVE      DOWN
```

```
Dell# show fcoe-map fid_1003
```

```
Fabric Name      fid_1003
Fabric Id        1003
Vlan Id          1003
Vlan priority    3
FC-MAP           0efc03
FKA-ADV-Period  8
Fcf Priority     128
Config-State     ACTIVE
Oper-State       UP
Members
Fc 0/0
Te 0/14 Te 0/16
```

Table 86. show fcoe-map Field Descriptions

Field	Description
Fabric-Name	Name of a SAN fabric.
Fabric ID	The ID number of the SAN fabric to which FC traffic is forwarded.
VLAN ID	The dedicated VLAN used to transport FCoE storage traffic between servers and a fabric over the NPG. The configured VLAN ID must be the same as the fabric ID.
VLAN priority	FCoE traffic uses VLAN priority 3. (This setting is not user-configurable.)
FC-MAP	FCoE MAC-address prefix value - The unique 24-bit MAC address prefix that identifies a fabric.
FKA-ADV-period	Time interval (in seconds) used to transmit FIP keepalive advertisements.
FCF Priority	The priority used by a server to select an upstream FCoE forwarder.
Config-State	Indicates whether the configured FCoE and FC parameters in the FCoE map are valid: Active (all mandatory FCoE and FC parameters are correctly configured) or Incomplete (either the FC-MAP value, fabric ID, or VLAN ID are not correctly configured).
Oper-State	Operational status of the link to the fabric: up (link is up and transmitting FC traffic), down (link is down and not transmitting FC traffic), link-wait (link is up and waiting for FLOGI to complete on peer FC port), or removed (port has been shut down).
Members	MXL 10/40GbE Switch or M I/O Aggregator with the FC Flex IO module Ethernet and FC ports, which are members of the dedicated FCoE VLAN that carries storage traffic to the specified fabric.

show qos dcb-map Command Examples

```
Dell# show qos dcb-map dcbmap2
```

```
State      :Complete  
PfcMode:ON
```

```
-----  
PG:0 TSA:ETS BW:50 PFC:OFF  
Priorities:0 1 2 4 5 6 7
```

```
PG:1 TSA:ETS BW:50 PFC:ON  
Priorities:3
```

Table 87. show qos dcb-map Field Descriptions

Field	Description
State	Complete: All mandatory DCB parameters are correctly configured. In progress: The DCB map configuration is not complete. Some mandatory parameters are not configured.
PFC Mode	PFC configuration in the DCB map: On (enabled) or Off.
PG	Priority group configured in the DCB map.
TSA	Transmission scheduling algorithm used in the DCB map: Enhanced Transmission Selection (ETS).
BW	Percentage of bandwidth allocated to the priority group.
PFC	PFC setting for the priority group: On (enabled) or Off.
Priorities	802.1p priorities configured in the priority group.

show npiv devices brief Command Example

```
Dell# show npiv devices brief
```

```
Total NPIV Devices = 2
```

```
-----
```

ENode-Intf LoginMethod	ENode-WWPN Status	FCoE-Vlan	Fabric-Intf	Fabric-Map
Te 0/12	20:01:00:10:18:f1:94:20	1003	Fc 0/5	
fid_1003	FLOGI	LOGGED_IN		
Te 0/13	10:00:00:00:c9:d9:9c:c0	1003	Fc 0/0	
fid_1003	FDISC	LOGGED_IN		

```
-----
```

Table 88. show npiv devices brief Field Descriptions

Field	Description
Total NPIV Devices	Number of downstream ENodes connected to a fabric over the MXL 10/40GbE Switch or M I/O Aggregator with the FC Flex IO module, NPG.
ENode-Intf	MXL 10/40GbE Switch or M I/O Aggregator with the FC Flex IO module Ethernet interface (<i>slot/port</i>) to which a server CNA is connected.
ENode-WWPN	Worldwide port name (WWPN) of a server CNA port.
FCoE-Vlan	VLAN ID of the dedicated VLAN used to transmit FCoE traffic to and from the fabric.
Fabric-Intf	Fabric-facing Fibre Channel port (<i>slot/port</i>) on which FC traffic is transmitted to the specified fabric.

Field	Description
Fabric-Map	Name of the FCoE map containing the FCoE/FC configuration parameters for the server CNA-fabric connection.
Login Method	Method used by the server CNA to log in to the fabric; for example: FLOGI - ENode logged in using a fabric login (FLOGI). FDISC - ENode logged in using a fabric discovery (FDISC).
Status	Operational status of the link between a server CNA port and a SAN fabric: Logged In - Server has logged in to the fabric and is able to transmit FCoE traffic

show npiv devices Command Example

```
Dell# show npiv devices
ENode[0]:
ENode MAC      :      00:10:18:f1:94:21
ENode Intf     :      Te 0/12
FCF MAC        :      5c:f9:dd:ef:10:c8
Fabric Intf    :      Fc 0/5
FCoE Vlan      :      1003
Fabric Map     :      fid_1003
ENode WWPN     :      20:01:00:10:18:f1:94:20
ENode WWNN     :      20:00:00:10:18:f1:94:21
FCoE MAC       :      0e:fc:03:01:02:01
FC-ID          :      01:02:01
LoginMethod    :      FLOGI
Secs           :      5593
Status         :      LOGGED_IN

ENode[1]:
ENode MAC      :      00:10:18:f1:94:22
ENode Intf     :      Te 0/13
FCF MAC        :      5c:f9:dd:ef:10:c9
Fabric Intf    :      Fc 0/0
FCoE Vlan      :      1003
Fabric Map     :      fid_1003
ENode WWPN     :      10:00:00:00:c9:d9:9c:cb
ENode WWNN     :      10:00:00:00:c9:d9:9c:cd
FCoE MAC       :      0e:fc:03:01:02:02
FC-ID          :      01:02:01
LoginMethod    :      FDISC
Secs           :      5593
Status         :      LOGGED_IN
```

Table 89. show npiv devices Field Descriptions

Field	Description
ENode <i>[number]</i>	Server CNA that has successfully logged in to a fabric over an MXL 10/40GbE Switch or M I/O Aggregator with the FC Flex IO module Ethernet port in ENode mode.
Enode MAC	MAC address of a server CNA port.
Enode Intf	Port number of a server-facing Ethernet port operating in ENode mode.

Field	Description
FCF MAC	Fibre Channel forwarder MAC: MAC address of MXL 10/40GbE Switch or M I/O Aggregator with the FC Flex IO module FCF interface.
Fabric Intf	Fabric-facing MXL 10/40GbE Switch or M I/O Aggregator with the FC Flex IO module Fibre Channel port (<i>slot/port</i>) on which FCoE traffic is transmitted to the specified fabric.
FCoE VLAN	ID of the dedicated VLAN used to transmit FCoE traffic from a server CNA to a fabric and configured on both the server-facing MXL 10/40GbE Switch or M I/O Aggregator with the FC Flex IO module port and server CNA port.
Fabric Map	Name of the FCoE map containing the FCoE/FC configuration parameters for the server CNA-fabric connection.
Enode WWPN	Worldwide port name of the server CNA port.
Enode WWNN	Worldwide node name of the server CNA.
FCoE MAC	Fabric-provided MAC address (FPMA). The FPMA consists of the FC-MAP value in the FCoE map and the FC-ID provided by the fabric after a successful FLOGI. In the FPMA, the most significant bytes are the FC-MAP; the least significant bytes are the FC-ID.
FC-ID	FC port ID provided by the fabric.
LoginMethod	Method used by the server CNA to log in to the fabric; for example, FLOGI or FDISC.
Secs	Number of seconds that the fabric connection is up.
State	Status of the fabric connection: logged in.

show fc switch Command Example

```
Dell# show fc switch
Switch Mode : NPG
Switch WWN  : 10:00:5c:f9:dd:ef:10:c0
Dell#
```

Table 90. show fc switch Command Description

Field	Description
Switch Mode	Fibre Channel mode of operation of an MXL 10/40GbE Switch or M I/O Aggregator with the FC Flex IO module. Default: NPG (configured as an NPIV proxy gateway).
Switch WWN	Factory-assigned worldwide node (WWN) name of the MXL 10/40GbE Switch or M I/O Aggregator with the FC Flex IO module. The MXL 10/40GbE Switch or M I/O Aggregator with the FC Flex IO module WWN name is not user-configurable.